

Amazon-Web-Services

Exam Questions SAA-C03

AWS Certified Solutions Architect - Associate (SAA-C03)



NEW QUESTION 1

- (Topic 1)

A company runs an online marketplace web application on AWS. The application serves hundreds of thousands of users during peak hours. The company needs a scalable, near-real-time solution to share the details of millions of financial transactions with several other internal applications. Transactions also need to be processed to remove sensitive data before being stored in a document database for low-latency retrieval.

What should a solutions architect recommend to meet these requirements?

- A. Store the transactions data into Amazon DynamoDB. Set up a rule in DynamoDB to remove sensitive data from every transaction upon write. Use DynamoDB Streams to share the transactions data with other applications.
- B. Stream the transactions data into Amazon Kinesis Data Firehose to store data in Amazon DynamoDB and Amazon S3. Use AWS Lambda integration with Kinesis Data Firehose to remove sensitive data.
- C. Other applications can consume the data stored in Amazon S3.
- D. Stream the transactions data into Amazon Kinesis Data Streams. Use AWS Lambda integration to remove sensitive data from every transaction and then store the transactions data in Amazon DynamoDB. Other applications can consume the transactions data off the Kinesis data stream.
- E. Store the batched transactions data in Amazon S3 as files.
- F. Use AWS Lambda to process every file and remove sensitive data before updating the files in Amazon S3. The Lambda function then stores the data in Amazon DynamoDB. Other applications can consume transaction files stored in Amazon S3.

Answer: C

Explanation:

The destination of your Kinesis Data Firehose delivery stream. Kinesis Data Firehose can send data records to various destinations, including Amazon Simple Storage Service (Amazon S3), Amazon Redshift, Amazon OpenSearch Service, and any HTTP endpoint that is owned by you or any of your third-party service providers. The following are the supported destinations:

- * Amazon OpenSearch Service
- * Amazon S3
- * Datadog
- * Dynatrace
- * Honeycomb
- * HTTP Endpoint
- * Logic Monitor
- * MongoDB Cloud
- * New Relic
- * Splunk
- * Sumo Logic <https://docs.aws.amazon.com/firehose/latest/dev/create-name.html>

<https://aws.amazon.com/kinesis/data-streams/>

Amazon Kinesis Data Streams (KDS) is a massively scalable and durable real-time data streaming service. KDS can continuously capture gigabytes of data per second from hundreds of thousands of sources such as website clickstreams, database event streams, financial transactions, social media feeds, IT logs, and location-tracking events.

NEW QUESTION 2

- (Topic 1)

A company hosts its web applications in the AWS Cloud. The company configures Elastic Load Balancers to use certificates that are imported into AWS Certificate Manager (ACM). The company's security team must be notified 30 days before the expiration of each certificate.

What should a solutions architect recommend to meet the requirement?

- A. Add a rule in ACM to publish a custom message to an Amazon Simple Notification Service (Amazon SNS) topic every day beginning 30 days before any certificate will expire.
- B. Create an AWS Config rule that checks for certificates that will expire within 30 days.
- C. Configure Amazon EventBridge (Amazon CloudWatch Events) to invoke a custom alert by way of Amazon Simple Notification Service (Amazon SNS) when AWS Config reports a noncompliant resource.
- D. Use AWS Trusted Advisor to check for certificates that will expire within 30 days.
- E. Create an Amazon CloudWatch alarm that is based on Trusted Advisor metrics for check status changes. Configure the alarm to send a custom alert by way of Amazon Simple Notification Service (Amazon SNS).
- F. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to detect any certificates that will expire within 30 days.
- G. Configure the rule to invoke an AWS Lambda function.
- H. Configure the Lambda function to send a custom alert by way of Amazon Simple Notification Service (Amazon SNS).

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/acm-certificate-expiration/>

NEW QUESTION 3

- (Topic 1)

A company has an application that provides marketing services to stores. The services are based on previous purchases by store customers. The stores upload transaction data to the company through SFTP, and the data is processed and analyzed to generate new marketing offers. Some of the files can exceed 200 GB in size.

Recently, the company discovered that some of the stores have uploaded files that contain personally identifiable information (PII) that should not have been included. The company wants administrators to be alerted if PII is shared again. The company also wants to automate remediation.

What should a solutions architect do to meet these requirements with the LEAST development effort?

- A. Use an Amazon S3 bucket as a secure transfer point.
- B. Use Amazon Inspector to scan the objects in the bucket.
- C. If objects contain PII.
- D. Trigger an S3 Lifecycle policy to remove the objects that contain PII.
- E. Use an Amazon S3 bucket as a secure transfer point.
- F. Use Amazon Macie to scan the objects in the bucket.
- G. If objects contain PII.

- H. Use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects that contain PII.
- I. Implement custom scanning algorithms in an AWS Lambda function
- J. Trigger the function when objects are loaded into the bucket
- K. If objects contain PII
- L. Use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects that contain PII.
- M. Implement custom scanning algorithms in an AWS Lambda function
- N. Trigger the function when objects are loaded into the bucket
- O. If objects contain PII
- P. Use Amazon Simple Email Service (Amazon SES) to trigger a notification to the administrators and trigger on S3 Lifecycle policy to remove the objects that contain PII.

Answer: B

Explanation:

To meet the requirements of detecting and alerting the administrators when PII is shared and automating remediation with the least development effort, the best approach would be to use Amazon S3 bucket as a secure transfer point and scan the objects in the bucket with Amazon Macie. Amazon Macie is a fully managed data security and data privacy service that uses machine learning and pattern matching to discover and protect sensitive data stored in Amazon S3. It can be used to classify sensitive data, monitor access to sensitive data, and automate remediation actions.

In this scenario, after uploading the files to the Amazon S3 bucket, the objects can be scanned for PII by Amazon Macie, and if it detects any PII, it can trigger an Amazon Simple Notification Service (SNS) notification to alert the administrators to remove the objects containing PII. This approach requires the least development effort, as Amazon Macie already has pre-built data classification rules that can detect PII in various formats. Hence, option B is the correct answer.

References:

? Amazon Macie User Guide: <https://docs.aws.amazon.com/macie/latest/userguide/what-is-macie.html>

? AWS Well-Architected Framework - Security Pillar: <https://docs.aws.amazon.com/wellarchitected/latest/security-pillar/welcome.html>

NEW QUESTION 4

- (Topic 1)

A company hosts more than 300 global websites and applications. The company requires a platform to analyze more than 30 TB of clickstream data each day. What should a solutions architect do to transmit and process the clickstream data?

- A. Design an AWS Data Pipeline to archive the data to an Amazon S3 bucket and run an Amazon EMR cluster with the data to generate analytics
- B. Create an Auto Scaling group of Amazon EC2 instances to process the data and send it to an Amazon S3 data lake for Amazon Redshift to use for analysis
- C. Cache the data to Amazon CloudFront. Store the data in an Amazon S3 bucket. When an object is added to the S3 bucket, run an AWS Lambda function to process the data for analysis.
- D. Collect the data from Amazon Kinesis Data Stream
- E. Use Amazon Kinesis Data Firehose to transmit the data to an Amazon S3 data lake. Load the data in Amazon Redshift for analysis

Answer: D

Explanation:

<https://aws.amazon.com/es/blogs/big-data/real-time-analytics-with-amazon-redshift-streaming-ingestion/>

NEW QUESTION 5

- (Topic 1)

A company runs a highly available image-processing application on Amazon EC2 instances in a single VPC. The EC2 instances run inside several subnets across multiple Availability Zones. The EC2 instances do not communicate with each other. However, the EC2 instances download images from Amazon S3 and upload images to Amazon S3 through a single NAT gateway. The company is concerned about data transfer charges.

What is the MOST cost-effective way for the company to avoid Regional data transfer charges?

- A. Launch the NAT gateway in each Availability Zone
- B. Replace the NAT gateway with a NAT instance
- C. Deploy a gateway VPC endpoint for Amazon S3
- D. Provision an EC2 Dedicated Host to run the EC2 instances

Answer: A

Explanation:

In this scenario, the company wants to avoid regional data transfer charges while downloading and uploading images from Amazon S3. To accomplish this at the lowest cost, the NAT gateway should be launched in each availability zone that the EC2 instances are running in. This allows the EC2 instances to route traffic through the local NAT gateway instead of sending traffic across an availability zone boundary and incurring regional data transfer fees. This method will help reduce the data transfer costs since inter- Availability Zone data transfers in a single region are free of charge.

Reference:

AWS NAT Gateway documentation: <https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html>

NEW QUESTION 6

- (Topic 1)

A company's website uses an Amazon EC2 instance store for its catalog of items. The company wants to make sure that the catalog is highly available and that the catalog is stored in a durable location.

What should a solutions architect do to meet these requirements?

- A. Move the catalog to Amazon ElastiCache for Redis.
- B. Deploy a larger EC2 instance with a larger instance store.
- C. Move the catalog from the instance store to Amazon S3 Glacier Deep Archive.
- D. Move the catalog to an Amazon Elastic File System (Amazon EFS) file system.

Answer: D

Explanation:

Moving the catalog to an Amazon Elastic File System (Amazon EFS) file system provides both high availability and durability. Amazon EFS is a fully-managed, highly-available, and durable file system that is built to scale on demand. With Amazon EFS, the catalog data can be stored and accessed from multiple EC2

instances in different availability zones, ensuring high availability. Also, Amazon EFS automatically stores files redundantly within and across multiple availability zones, making it a durable storage option.

NEW QUESTION 7

- (Topic 1)

A company has a three-tier web application that is deployed on AWS. The web servers are deployed in a public subnet in a VPC. The application servers and database servers are deployed in private subnets in the same VPC. The company has deployed a third-party virtual firewall appliance from AWS Marketplace in an inspection VPC. The appliance is configured with an IP interface that can accept IP packets. A solutions architect needs to integrate the web application with the appliance to inspect all traffic to the application before the traffic reaches the web server. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create a Network Load Balancer in the public subnet of the application's VPC to route the traffic to the appliance for packet inspection
- B. Create an Application Load Balancer in the public subnet of the application's VPC to route the traffic to the appliance for packet inspection
- C. Deploy a transit gateway in the inspection VPC. Configure route tables to route the incoming packets through the transit gateway
- D. Deploy a Gateway Load Balancer in the inspection VPC. Create a Gateway Load Balancer endpoint to receive the incoming packets and forward the packets to the appliance

Answer: D

Explanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/scaling-network-traffic-inspection-using-aws-gateway-load-balancer/>

NEW QUESTION 8

- (Topic 1)

A company is migrating a distributed application to AWS. The application serves variable workloads. The legacy platform consists of a primary server that coordinates jobs across multiple compute nodes. The company wants to modernize the application with a solution that maximizes resiliency and scalability. How should a solutions architect design the architecture to meet these requirements?

- A. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination for the jobs. Implement the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group
- B. Configure EC2 Auto Scaling to use scheduled scaling
- C. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination for the jobs. Implement the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group. Configure EC2 Auto Scaling based on the size of the queue
- D. Implement the primary server and the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group
- E. Configure AWS CloudTrail as a destination for the jobs. Configure EC2 Auto Scaling based on the load on the primary server
- F. Implement the primary server and the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group. Configure Amazon EventBridge (Amazon CloudWatch Events) as a destination for the jobs. Configure EC2 Auto Scaling based on the load on the compute nodes

Answer: B

Explanation:

To maximize resiliency and scalability, the best solution is to use an Amazon SQS queue as a destination for the jobs. This decouples the primary server from the compute nodes, allowing them to scale independently. This also helps to prevent job loss in the event of a failure. Using an Auto Scaling group of Amazon EC2 instances for the compute nodes allows for automatic scaling based on the workload. In this case, it's recommended to configure the Auto Scaling group based on the size of the Amazon SQS queue, which is a better indicator of the actual workload than the load on the primary server or compute nodes. This approach ensures that the application can handle variable workloads, while also minimizing costs by automatically scaling up or down the compute nodes as needed.

NEW QUESTION 9

- (Topic 1)

A company has a production workload that runs on 1,000 Amazon EC2 Linux instances. The workload is powered by third-party software. The company needs to patch the third-party software on all EC2 instances as quickly as possible to remediate a critical security vulnerability. What should a solutions architect do to meet these requirements?

- A. Create an AWS Lambda function to apply the patch to all EC2 instances.
- B. Configure AWS Systems Manager Patch Manager to apply the patch to all EC2 instances.
- C. Schedule an AWS Systems Manager maintenance window to apply the patch to all EC2 instances.
- D. Use AWS Systems Manager Run Command to run a custom command that applies the patch to all EC2 instances.

Answer: B

Explanation:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/about-windows-app-patching.html>

NEW QUESTION 10

- (Topic 1)

A company is developing an application that provides order shipping statistics for retrieval by a REST API. The company wants to extract the shipping statistics, organize the data into an easy-to-read HTML format, and send the report to several email addresses at the same time every morning. Which combination of steps should a solutions architect take to meet these requirements? (Choose two.)

- A. Configure the application to send the data to Amazon Kinesis Data Firehose.
- B. Use Amazon Simple Email Service (Amazon SES) to format the data and to send the report by email.
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled event that invokes an AWS Glue job to query the application's API for the data.
- D. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled event that invokes an AWS Lambda function to query the application's API for the data.
- E. Store the application data in Amazon S3. Create an Amazon Simple Notification Service (Amazon SNS) topic as an S3 event destination to send the report by

Answer: BD

Explanation:

<https://docs.aws.amazon.com/ses/latest/dg/send-email-formatted.html>

* D. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled event that invokes an AWS Lambda function to query the application's API for the data. This step can be done using AWS Lambda to extract the shipping statistics and organize the data into an HTML format.

* B. Use Amazon Simple Email Service (Amazon SES) to format the data and send the report by email. This step can be done by using Amazon SES to send the report to multiple email addresses at the same time every morning.

Therefore, options D and B are the correct choices for this question. Option A is incorrect because Kinesis Data Firehose is not necessary for this use case. Option C is incorrect because AWS Glue is not required to query the application's API. Option E is incorrect because S3 event notifications cannot be used to send the report by email.

NEW QUESTION 10

- (Topic 1)

A company is preparing to launch a public-facing web application in the AWS Cloud. The architecture consists of Amazon EC2 instances within a VPC behind an Elastic Load Balancer (ELB). A third-party service is used for the DNS. The company's solutions architect must recommend a solution to detect and protect against large-scale DDoS attacks.

Which solution meets these requirements?

- A. Enable Amazon GuardDuty on the account.
- B. Enable Amazon Inspector on the EC2 instances.
- C. Enable AWS Shield and assign Amazon Route 53 to it.
- D. Enable AWS Shield Advanced and assign the ELB to it.

Answer: D

Explanation:

<https://aws.amazon.com/shield/faqs/>

NEW QUESTION 15

- (Topic 1)

A company's HTTP application is behind a Network Load Balancer (NLB). The NLB's target group is configured to use an Amazon EC2 Auto Scaling group with multiple EC2 instances that run the web service.

The company notices that the NLB is not detecting HTTP errors for the application. These errors require a manual restart of the EC2 instances that run the web service. The company needs to improve the application's availability without writing custom scripts or code.

What should a solutions architect do to meet these requirements?

- A. Enable HTTP health checks on the NL
- B. supplying the URL of the company's application.
- C. Add a cron job to the EC2 instances to check the local application's logs once each minut
- D. If HTTP errors are detected, the application will restart.
- E. Replace the NLB with an Application Load Balance
- F. Enable HTTP health checks by supplying the URL of the company's applicatio
- G. Configure an Auto Scaling action to replace unhealthy instances.
- H. Create an Amazon Cloud Watch alarm that monitors the UnhealthyHostCount metric for the NL
- I. Configure an Auto Scaling action to replace unhealthy instances when the alarm is in the ALARM state.

Answer: C

Explanation:

Application availability: NLB cannot assure the availability of the application. This is because it bases its decisions solely on network and TCP-layer variables and has no awareness of the application at all. Generally, NLB determines availability based on the ability of a server to respond to ICMP ping or to correctly complete the three-way TCP handshake. ALB goes much deeper and is capable of determining availability based on not only a successful HTTP GET of a particular page but also the verification that the content is as was expected based on the input parameters.

NEW QUESTION 19

- (Topic 1)

A company is hosting a static website on Amazon S3 and is using Amazon Route 53 for DNS. The website is experiencing increased demand from around the world. The company must decrease latency for users who access the website.

Which solution meets these requirements MOST cost-effectively?

- A. Replicate the S3 bucket that contains the website to all AWS Region
- B. Add Route 53 geolocation routing entries.
- C. Provision accelerators in AWS Global Accelerato
- D. Associate the supplied IP addresses with the S3 bucke
- E. Edit the Route 53 entries to point to the IP addresses of the accelerators.
- F. Add an Amazon CloudFront distribution in front of the S3 bucke
- G. Edit the Route 53 entries to point to the CloudFront distribution.
- H. Enable S3 Transfer Acceleration on the bucke
- I. Edit the Route 53 entries to point to the new endpoint.

Answer: C

Explanation:

Amazon CloudFront is a content delivery network (CDN) that caches content at edge locations around the world, providing low latency and high transfer speeds to users accessing the content. Adding a CloudFront distribution in front of the S3 bucket will cache the static website's content at edge locations around the world, decreasing latency for users accessing the website. This solution is also cost-effective as it only charges for the data transfer and requests made by users accessing the content from the CloudFront edge locations. Additionally, this solution provides scalability and reliability benefits as CloudFront can automatically scale to handle increased demand and provide high availability for the website.

NEW QUESTION 24

- (Topic 1)

A company hosts an application on AWS Lambda functions that are invoked by an Amazon API Gateway API. The Lambda functions save customer data to an Amazon Aurora MySQL database. Whenever the company upgrades the database, the Lambda functions fail to establish database connections until the upgrade is complete. The result is that customer data is not recorded for some of the event.

A solutions architect needs to design a solution that stores customer data that is created during database upgrades. Which solution will meet these requirements?

- A. Provision an Amazon RDS proxy to sit between the Lambda functions and the database. Configure the Lambda functions to connect to the RDS proxy.
- B. Increase the run time of the Lambda functions to the maximum. Create a retry mechanism in the code that stores the customer data in the database.
- C. Persist the customer data to Lambda local storage.
- D. Configure new Lambda functions to scan the local storage to save the customer data to the database.
- E. Store the customer data in an Amazon Simple Queue Service (Amazon SQS) FIFO queue. Create a new Lambda function that polls the queue and stores the customer data in the database.

Answer: D

Explanation:

<https://www.learnaws.org/2020/12/13/aws-rds-proxy-deep-dive/>

RDS proxy can improve application availability in such a situation by waiting for the new database instance to be functional and maintaining any requests received from the application during this time. The end result is that the application is more resilient to issues with the underlying database.

This will enable solution to hold data till the time DB comes back to normal. RDS proxy is to optimally utilize the connection between Lambda and DB. Lambda can open multiple connection concurrently which can be taxing on DB compute resources, hence RDS proxy was introduced to manage and leverage these connections efficiently.

NEW QUESTION 28

- (Topic 1)

A company is launching a new application and will display application metrics on an Amazon CloudWatch dashboard. The company's product manager needs to access this dashboard periodically. The product manager does not have an AWS account. A solution architect must provide access to the product manager by following the principle of least privilege.

Which solution will meet these requirements?

- A. Share the dashboard from the CloudWatch console.
- B. Enter the product manager's email address, and complete the sharing step.
- C. Provide a shareable link for the dashboard to the product manager.
- D. Create an IAM user specifically for the product manager.
- E. Attach the CloudWatch Read Only Access managed policy to the user.
- F. Share the new login credential with the product manager.
- G. Share the browser URL of the correct dashboard with the product manager.
- H. Create an IAM user for the company's employees, Attach the View Only Access AWS managed policy to the IAM user.
- I. Share the new login credentials with the product manager.
- J. Ask the product manager to navigate to the CloudWatch console and locate the dashboard by name in the Dashboards section.
- K. Deploy a bastion server in a public subnet.
- L. When the product manager requires access to the dashboard, start the server and share the RDP credential.
- M. On the bastion server, ensure that the browser is configured to open the dashboard URL with cached AWS credentials that have appropriate permissions to view the dashboard.

Answer: B

Explanation:

To provide the product manager access to the Amazon CloudWatch dashboard while following the principle of least privilege, a solution architect should create an IAM user specifically for the product manager and attach the CloudWatch Read Only Access managed policy to the user. This policy allows the user to view the dashboard without being able to make any changes to it. The solution architect should then share the new login credential with the product manager and provide them with the browser URL of the correct dashboard.

NEW QUESTION 33

- (Topic 1)

A company wants to migrate an on-premises data center to AWS. The data center hosts an SFTP server that stores its data on an NFS-based file system. The server holds 200 GB of data that needs to be transferred. The server must be hosted on an Amazon EC2 instance that uses an Amazon Elastic File System (Amazon EFS) file system.

When combination of steps should a solutions architect take to automate this task? (Select TWO)

- A. Launch the EC2 instance into the same Availability Zone as the EFS file system.
- B. Install an AWS DataSync agent in the on-premises data center.
- C. Create a secondary Amazon Elastic Block Store (Amazon EBS) volume on the EC2 instance for the data.
- D. Manually use an operating system copy command to push the data to the EC2 instance.
- E. Use AWS DataSync to create a suitable location configuration for the on-premises SFTP server.

Answer: BE

Explanation:

AWS DataSync is an online data movement and discovery service that simplifies data migration and helps users quickly, easily, and securely move their file or object data to, from, and between AWS storage services¹. Users can use AWS DataSync to transfer data between on-premises and AWS storage services. To use AWS DataSync, users need to install an AWS DataSync agent in the on-premises data center. The agent is a software appliance that connects to the source or destination storage system and handles the data transfer to or from AWS over the network². Users also need to use AWS DataSync to create a suitable location configuration for the on-premises SFTP server. A location is a logical representation of a storage system that contains files or objects that users want to transfer using DataSync. Users can create locations for NFS shares, SMB shares, HDFS file systems, self-managed object storage, Amazon S3 buckets, Amazon EFS file systems, Amazon FSx for Windows File Server file systems, Amazon FSx for Lustre file systems, Amazon FSx for OpenZFS file systems, Amazon FSx for NetApp ONTAP file systems, and AWS Snowcone devices³.

NEW QUESTION 37

- (Topic 1)

A company needs to keep user transaction data in an Amazon DynamoDB table. The company must retain the data for 7 years. What is the MOST operationally efficient solution that meets these requirements?

- A. Use DynamoDB point-in-time recovery to back up the table continuously.
- B. Use AWS Backup to create backup schedules and retention policies for the table.
- C. Create an on-demand backup of the table by using the DynamoDB console.
- D. Store the backup in an Amazon S3 bucket.
- E. Set an S3 Lifecycle configuration for the S3 bucket.
- F. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function.
- G. Configure the Lambda function to back up the table and to store the backup in an Amazon S3 bucket.
- H. Set an S3 Lifecycle configuration for the S3 bucket.

Answer: C

NEW QUESTION 41

- (Topic 1)

An ecommerce company wants to launch a one-deal-a-day website on AWS. Each day will feature exactly one product on sale for a period of 24 hours. The company wants to be able to handle millions of requests each hour with millisecond latency during peak hours. Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon S3 to host the full website in different S3 buckets. Add Amazon CloudFront distributions. Set the S3 buckets as origins for the distributions. Store the order data in Amazon S3.
- B. Deploy the full website on Amazon EC2 instances that run in Auto Scaling groups across multiple Availability Zones. Add an Application Load Balancer (ALB) to distribute the website traffic. Add another ALB for the backend APIs. Store the data in Amazon RDS for MySQL.
- C. Migrate the full application to run in containers. Host the containers on Amazon Elastic Kubernetes Service (Amazon EKS). Use the Kubernetes Cluster Autoscaler to increase and decrease the number of pods to process bursts in traffic. Store the data in Amazon RDS for MySQL.
- D. Use an Amazon S3 bucket to host the website's static content. Deploy an Amazon CloudFront distribution.
- E. Set the S3 bucket as the origin. Use Amazon API Gateway and AWS Lambda functions for the backend APIs. Store the data in Amazon DynamoDB.

Answer: D

Explanation:

To launch a one-deal-a-day website on AWS with millisecond latency during peak hours and with the least operational overhead, the best option is to use an Amazon S3 bucket to host the website's static content, deploy an Amazon CloudFront distribution, set the S3 bucket as the origin, use Amazon API Gateway and AWS Lambda functions for the backend APIs, and store the data in Amazon DynamoDB. This option requires minimal operational overhead and can handle millions of requests each hour with millisecond latency during peak hours. Therefore, option D is the correct answer.

Reference: <https://aws.amazon.com/blogs/compute/building-a-serverless-multi-player-game-with-aws-lambda-and-amazon-dynamodb/>

NEW QUESTION 42

- (Topic 1)

A company is implementing a shared storage solution for a media application that is hosted in the AWS Cloud. The company needs the ability to use SMB clients to access data. The solution must be fully managed. Which AWS solution meets these requirements?

- A. Create an AWS Storage Gateway volume gateway.
- B. Create a file share that uses the required client protocol. Connect the application server to the file share.
- C. Create an AWS Storage Gateway tape gateway. Configure (apes) to use Amazon S3. Connect the application server to the tape gateway.
- D. Create an Amazon EC2 Windows instance. Install and configure a Windows file share role on the instance.
- E. Connect the application server to the file share.
- F. Create an Amazon FSx for Windows File Server file system. Attach the file system to the origin server.
- G. Connect the application server to the file system.

Answer: D

Explanation:

<https://aws.amazon.com/fsx/lustre/>

Amazon FSx has native support for Windows file system features and for the industry-standard Server Message Block (SMB) protocol to access file storage over a network. <https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what-is.html>

NEW QUESTION 47

- (Topic 1)

A company has registered its domain name with Amazon Route 53. The company uses Amazon API Gateway in the ca-central-1 Region as a public interface for its backend microservice APIs. Third-party services consume the APIs securely. The company wants to design its API Gateway URL with the company's domain name and corresponding certificate so that the third-party services can use HTTPS.

Which solution will meet these requirements?

- A. Create stage variables in API Gateway with Name="Endpoint-URL" and Value="Company Domain Name" to overwrite the default URL.
- B. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM).
- C. Create Route 53 DNS records with the company's domain name.
- D. Point the alias record to the Regional API Gateway stage endpoint.
- E. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM) in the us-east-1 Region.
- F. Create a Regional API Gateway endpoint.
- G. Associate the API Gateway endpoint with the company's domain name.
- H. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM) in the same Region.
- I. Attach the certificate to the API Gateway endpoint.
- J. Configure Route 53 to route traffic to the API Gateway endpoint.
- K. Create a Regional API Gateway endpoint.
- L. Associate the API Gateway endpoint with the company's domain name.
- M. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM) in the us-east-1 Region.
- N. Attach the certificate to the API Gateway API.

- O. Create Route 53 DNS records with the company's domain name.
- P. Point an A record to the company's domain name.

Answer: C

Explanation:

To design the API Gateway URL with the company's domain name and corresponding certificate, the company needs to do the following: 1. Create a Regional API Gateway endpoint: This will allow the company to create an endpoint that is specific to a region. 2. Associate the API Gateway endpoint with the company's domain name: This will allow the company to use its own domain name for the API Gateway URL. 3. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM) in the same Region: This will allow the company to use HTTPS for secure communication with its APIs. 4. Attach the certificate to the API Gateway endpoint: This will allow the company to use the certificate for securing the API Gateway URL. 5. Configure Route 53 to route traffic to the API Gateway endpoint: This will allow the company to use Route 53 to route traffic to the API Gateway URL using the company's domain name.

NEW QUESTION 48

- (Topic 1)

A company is running an SMB file server in its data center. The file server stores large files that are accessed frequently for the first few days after the files are created. After 7 days the files are rarely accessed.

The total data size is increasing and is close to the company's total storage capacity. A solutions architect must increase the company's available storage space without losing low-latency access to the most recently accessed files. The solutions architect must also provide file lifecycle management to avoid future storage issues.

Which solution will meet these requirements?

- A. Use AWS DataSync to copy data that is older than 7 days from the SMB file server to AWS.
- B. Create an Amazon S3 File Gateway to extend the company's storage space.
- C. Create an S3 Lifecycle policy to transition the data to S3 Glacier Deep Archive after 7 days.
- D. Create an Amazon FSx for Windows File Server file system to extend the company's storage space.
- E. Install a utility on each user's computer to access Amazon S3. Create an S3 Lifecycle policy to transition the data to S3 Glacier Flexible Retrieval after 7 days.

Answer: B

Explanation:

Amazon S3 File Gateway is a hybrid cloud storage service that enables on-premises applications to seamlessly use Amazon S3 cloud storage. It provides a file interface to Amazon S3 and supports SMB and NFS protocols. It also supports S3 Lifecycle policies that can automatically transition data from S3 Standard to S3 Glacier Deep Archive after a specified period of time. This solution will meet the requirements of increasing the company's available storage space without losing low-latency access to the most recently accessed files and providing file lifecycle management to avoid future storage issues.

Reference:

<https://docs.aws.amazon.com/storagegateway/latest/userguide/WhatIsStorageGateway.html>

NEW QUESTION 52

- (Topic 1)

A company hosts an application on multiple Amazon EC2 instances. The application processes messages from an Amazon SQS queue, writes to an Amazon RDS table, and deletes the message from the queue. Occasional duplicate records are found in the RDS table. The SQS queue does not contain any duplicate messages.

What should a solutions architect do to ensure messages are being processed once only?

- A. Use the CreateQueue API call to create a new queue.
- B. Use the AddPermission API call to add appropriate permissions.
- C. Use the ReceiveMessage API call to set an appropriate wait time.
- D. Use the ChangeMessageVisibility API call to increase the visibility timeout.

Answer: D

Explanation:

The visibility timeout begins when Amazon SQS returns a message. During this time, the consumer processes and deletes the message. However, if the consumer fails before deleting the message and your system doesn't call the DeleteMessage action for that message before the visibility timeout expires, the message becomes visible to other consumers and the message is received again. If a message must be received only once, your consumer should delete it within the duration of the visibility timeout. <https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-visibility-timeout.html>

Keyword: SQS queue writes to an Amazon RDS. From this, Option D is the best choice; other options are ruled out [Option A - You can't introduce one more queue in the existing one; Option B - only Permission; Option C - Only Retrieves Messages]. FIFO queues are designed to never introduce duplicate messages. However, your message producer might introduce duplicates in certain scenarios: for example, if the producer sends a message, does not receive a response, and then resends the same message. Amazon SQS APIs provide deduplication functionality that prevents your message producer from sending duplicates. Any duplicates introduced by the message producer are removed within a 5-minute deduplication interval. For standard queues, you might occasionally receive a duplicate copy of a message (at-least- once delivery). If you use a standard queue, you must design your applications to be idempotent (that is, they must not be affected adversely when processing the same message more than once).

NEW QUESTION 53

- (Topic 1)

A company needs to review its AWS Cloud deployment to ensure that its Amazon S3 buckets do not have unauthorized configuration changes.

What should a solutions architect do to accomplish this goal?

- A. Turn on AWS Config with the appropriate rules.
- B. Turn on AWS Trusted Advisor with the appropriate checks.
- C. Turn on Amazon Inspector with the appropriate assessment template.
- D. Turn on Amazon S3 server access logging.
- E. Configure Amazon EventBridge (Amazon Cloud Watch Events).

Answer: A

Explanation:

To ensure that Amazon S3 buckets do not have unauthorized configuration changes, a solutions architect should turn on AWS Config with the appropriate rules.

AWS Config is a service that allows users to audit and assess their AWS resource configurations for compliance with industry standards and internal policies. It provides a detailed view of the resources and their configurations, including information on how the resources are related to each other. By turning on AWS Config with the appropriate rules, users can identify and remediate unauthorized configuration changes to their Amazon S3 buckets.

NEW QUESTION 56

- (Topic 1)

A company collects temperature, humidity, and atmospheric pressure data in cities across multiple continents. The average volume of data collected per site each day is 500 GB. Each site has a high-speed internet connection. The company's weather forecasting applications are based in a single Region and analyze the data daily.

What is the FASTEST way to aggregate data from all of these global sites?

- A. Enable Amazon S3 Transfer Acceleration on the destination bucket
- B. Use multipart uploads to directly upload site data to the destination bucket.
- C. Upload site data to an Amazon S3 bucket in the closest AWS Region
- D. Use S3 cross-Region replication to copy objects to the destination bucket.
- E. Schedule AWS Snowball jobs daily to transfer data to the closest AWS Region
- F. Use S3 cross-Region replication to copy objects to the destination bucket.
- G. Upload the data to an Amazon EC2 instance in the closest Region
- H. Store the data in an Amazon Elastic Block Store (Amazon EBS) volume
- I. Once a day take an EBS snapshot and copy it to the centralized Region
- J. Restore the EBS volume in the centralized Region and run an analysis on the data daily.

Answer: A

Explanation:

You might want to use Transfer Acceleration on a bucket for various reasons, including the following:

You have customers that upload to a centralized bucket from all over the world. You transfer gigabytes to terabytes of data on a regular basis across continents.

You are unable to utilize all of your available bandwidth over the Internet when uploading to Amazon S3.

<https://docs.aws.amazon.com/AmazonS3/latest/dev/transfer-acceleration.html>

[https://aws.amazon.com/s3/transfer-acceleration/#:~:text=S3%20Transfer%20Acceleration%20\(S3TA\)%20reduces,to%20S3%20for%20remote%20applications:](https://aws.amazon.com/s3/transfer-acceleration/#:~:text=S3%20Transfer%20Acceleration%20(S3TA)%20reduces,to%20S3%20for%20remote%20applications:)

"Amazon S3 Transfer Acceleration can speed up content transfers to and from Amazon S3 by as much as 50-500% for long-distance transfer of larger objects.

Customers who have either web or mobile applications with widespread users or applications hosted far away from their S3 bucket can experience long and variable upload and download speeds over the Internet"

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/mpuoverview.html>

"Improved throughput - You can upload parts in parallel to improve throughput."

NEW QUESTION 61

- (Topic 1)

A solutions architect must design a highly available infrastructure for a website. The website is powered by Windows web servers that run on Amazon EC2 instances. The solutions architect must implement a solution that can mitigate a large-scale DDoS attack that originates from thousands of IP addresses.

Downtime is not acceptable for the website.

Which actions should the solutions architect take to protect the website from such an attack? (Select TWO.)

- A. Use AWS Shield Advanced to stop the DDoS attack.
- B. Configure Amazon GuardDuty to automatically block the attackers.
- C. Configure the website to use Amazon CloudFront for both static and dynamic content.
- D. Use an AWS Lambda function to automatically add attacker IP addresses to VPC network ACLs.
- E. Use EC2 Spot Instances in an Auto Scaling group with a target tracking scaling policy that is set to 80% CPU utilization

Answer: AC

Explanation:

(<https://aws.amazon.com/cloudfront>)

NEW QUESTION 62

- (Topic 1)

A company recently signed a contract with an AWS Managed Service Provider (MSP) Partner for help with an application migration initiative. A solutions architect needs to share an Amazon Machine Image (AMI) from an existing AWS account with the MSP Partner's AWS account. The AMI is backed by Amazon Elastic Block Store (Amazon EBS) and uses a customer managed customer master key (CMK) to encrypt EBS volume snapshots.

The AMI is backed by Amazon Elastic Block Store (Amazon EBS) and uses a customer managed customer master key (CMK) to encrypt EBS volume snapshots.

What is the MOST secure way for the solutions architect to share the AMI with the MSP Partner's AWS account?

- A. Make the encrypted AMI and snapshots publicly available
- B. Modify the CMK's key policy to allow the MSP Partner's AWS account to use the key
- C. Modify the launchPermission property of the AMI
- D. Share the AMI with the MSP Partner's AWS account only
- E. Modify the CMK's key policy to allow the MSP Partner's AWS account to use the key.
- F. Modify the launchPermission property of the AMI. Share the AMI with the MSP Partner's AWS account only
- G. Modify the CMK's key policy to trust a new CMK that is owned by the MSP Partner for encryption.
- H. Export the AMI from the source account to an Amazon S3 bucket in the MSP Partner's AWS account
- I. Encrypt the S3 bucket with a CMK that is owned by the MSP Partner. Copy and launch the AMI in the MSP Partner's AWS account.

Answer: B

Explanation:

Share the existing KMS key with the MSP external account because it has already been used to encrypt the AMI snapshot.

<https://docs.aws.amazon.com/kms/latest/developerguide/key-policy-modifying-external-accounts.html>

NEW QUESTION 65

- (Topic 1)

A company runs a photo processing application that needs to frequently upload and download pictures from Amazon S3 buckets that are located in the same AWS Region. A solutions architect has noticed an increased cost in data transfer fees and needs to implement a solution to reduce these costs. How can the solutions architect meet this requirement?

- A. Deploy Amazon API Gateway into a public subnet and adjust the route table to route S3 calls through it.
- B. Deploy a NAT gateway into a public subnet and attach an end point policy that allows access to the S3 buckets.
- C. Deploy the application into a public subnet and allow it to route through an internet gateway to access the S3 Buckets
- D. Deploy an S3 VPC gateway endpoint into the VPC and attach an endpoint policy that allows access to the S3 buckets.

Answer: D

Explanation:

The correct answer is Option D. Deploy an S3 VPC gateway endpoint into the VPC and attach an endpoint policy that allows access to the S3 buckets. By deploying an S3 VPC gateway endpoint, the application can access the S3 buckets over a private network connection within the VPC, eliminating the need for data transfer over the internet. This can help reduce data transfer fees as well as improve the performance of the application. The endpoint policy can be used to specify which S3 buckets the application has access to.

NEW QUESTION 69

- (Topic 1)

A solutions architect is developing a multiple-subnet VPC architecture. The solution will consist of six subnets in two Availability Zones. The subnets are defined as public, private and dedicated for databases. Only the Amazon EC2 instances running in the private subnets should be able to access a database. Which solution meets these requirements?

- A. Create a new route table that excludes the route to the public subnets' CIDR block
- B. Associate the route table to the database subnets.
- C. Create a security group that denies ingress from the security group used by instances in the public subnet
- D. Attach the security group to an Amazon RDS DB instance.
- E. Create a security group that allows ingress from the security group used by instances in the private subnet
- F. Attach the security group to an Amazon RDS DB instance.
- G. Create a new peering connection between the public subnets and the private subnet
- H. Create a different peering connection between the private subnets and the databasesubnets.

Answer: C

Explanation:

Security groups are stateful. All inbound traffic is blocked by default. If you create an inbound rule allowing traffic in, that traffic is automatically allowed back out again. You cannot block specific IP address using Security groups (instead use Network Access Control Lists).

"You can specify allow rules, but not deny rules." "When you first create a security group, it has no inbound rules. Therefore, no inbound traffic originating from another host to your instance is allowed until you add inbound rules to the security group." Source:

https://docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html#VPCSecurityGroups

NEW QUESTION 70

- (Topic 1)

A company is migrating applications to AWS. The applications are deployed in different accounts. The company manages the accounts centrally by using AWS Organizations. The company's security team needs a single sign-on (SSO) solution across all the company's accounts. The company must continue managing the users and groups in its on-premises self-managed Microsoft Active Directory. Which solution will meet these requirements?

- A. Enable AWS Single Sign-On (AWS SSO) from the AWS SSO console
- B. Create a one-way forest trust or a one-way domain trust to connect the company's self-managed Microsoft Active Directory with AWS SSO by using AWS Directory Service for Microsoft Active Directory.
- C. Enable AWS Single Sign-On (AWS SSO) from the AWS SSO console
- D. Create a two-way forest trust to connect the company's self-managed Microsoft Active Directory with AWS SSO by using AWS Directory Service for Microsoft Active Directory.
- E. Use AWS Directory Service
- F. Create a two-way trust relationship with the company's self-managed Microsoft Active Directory.
- G. Deploy an identity provider (IdP) on premise
- H. Enable AWS Single Sign-On (AWS SSO) from the AWS SSO console.

Answer: A

Explanation:

To provide single sign-on (SSO) across all the company's accounts while continuing to manage users and groups in its on-premises self-managed Microsoft Active Directory, the solution is to enable AWS Single Sign-On (SSO) from the AWS SSO console and create a one-way forest trust or a one-way domain trust to connect the company's self-managed Microsoft Active Directory with AWS SSO by using AWS Directory Service for Microsoft Active Directory. This solution is described in the AWS documentation

NEW QUESTION 72

- (Topic 1)

A company uses Amazon S3 to store its confidential audit documents. The S3 bucket uses bucket policies to restrict access to audit team IAM user credentials according to the principle of least privilege. Company managers are worried about accidental deletion of documents in the S3 bucket and want a more secure solution. What should a solutions architect do to secure the audit documents?

- A. Enable the versioning and MFA Delete features on the S3 bucket.
- B. Enable multi-factor authentication (MFA) on the IAM user credentials for each auditteam IAM user account.
- C. Add an S3 Lifecycle policy to the audit team's IAM user accounts to deny the s3:DeleteObject action during audit dates.
- D. Use AWS Key Management Service (AWS KMS) to encrypt the S3 bucket and restrict audit team IAM user accounts from accessing the KMS key.

Answer: A

NEW QUESTION 77

- (Topic 1)

An application allows users at a company's headquarters to access product data. The product data is stored in an Amazon RDS MySQL DB instance. The operations team has isolated an application performance slowdown and wants to separate read traffic from write traffic. A solutions architect needs to optimize the application's performance quickly.

What should the solutions architect recommend?

- A. Change the existing database to a Multi-AZ deployment
- B. Serve the read requests from the primary Availability Zone.
- C. Change the existing database to a Multi-AZ deployment
- D. Serve the read requests from the secondary Availability Zone.
- E. Create read replicas for the database
- F. Configure the read replicas with half of the compute and storage resources as the source database.
- G. Create read replicas for the database
- H. Configure the read replicas with the same compute and storage resources as the source database.

Answer: D

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_MySQL.Replication.ReadReplicas.html

NEW QUESTION 82

- (Topic 1)

A solutions architect is designing a VPC with public and private subnets. The VPC and subnets use IPv4 CIDR blocks. There is one public subnet and one private subnet in each of three Availability Zones (AZs) for high availability. An internet gateway is used to provide internet access for the public subnets. The private subnets require access to the internet to allow Amazon EC2 instances to download software updates.

What should the solutions architect do to enable Internet access for the private subnets?

- A. Create three NAT gateways, one for each public subnet in each AZ
- B. Create a private route table for each AZ that forwards non-VPC traffic to the NAT gateway in its AZ.
- C. Create three NAT instances, one for each private subnet in each AZ
- D. Create a private route table for each AZ that forwards non-VPC traffic to the NAT instance in its AZ.
- E. Create a second internet gateway on one of the private subnets
- F. Update the route table for the private subnets that forward non-VPC traffic to the private internet gateway.
- G. Create an egress-only internet gateway on one of the public subnets
- H. Update the route table for the private subnets that forward non-VPC traffic to the egress-only internet gateway.

Answer: A

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2018/03/introducing-amazon-vpc-nat-gateway-in-the-aws-govcloud-us-region/#:~:text=NAT%20Gateway%20is%20a%20highly,instances%20in%20a%20private%20subnet.>

<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-comparison.html>

NEW QUESTION 87

- (Topic 1)

An application runs on an Amazon EC2 instance in a VPC. The application processes logs that are stored in an Amazon S3 bucket. The EC2 instance needs to access the S3 bucket without connectivity to the internet.

Which solution will provide private network connectivity to Amazon S3?

- A. Create a gateway VPC endpoint to the S3 bucket.
- B. Stream the logs to Amazon CloudWatch Log
- C. Export the logs to the S3 bucket.
- D. Create an instance profile on Amazon EC2 to allow S3 access.
- E. Create an Amazon API Gateway API with a private link to access the S3 endpoint.

Answer: A

Explanation:

VPC endpoint allows you to connect to AWS services using a private network instead of using the public Internet

NEW QUESTION 88

- (Topic 1)

A company has thousands of edge devices that collectively generate 1 TB of status alerts each day. Each alert is approximately 2 KB in size. A solutions architect needs to implement a solution to ingest and store the alerts for future analysis.

The company wants a highly available solution. However, the company needs to minimize costs and does not want to manage additional infrastructure. Additionally, the company wants to keep 14 days of data available for immediate analysis and archive any data older than 14 days.

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

- A. Create an Amazon Kinesis Data Firehose delivery stream to ingest the alerts Configure the Kinesis Data Firehose stream to deliver the alerts to an Amazon S3 bucket Set up an S3 Lifecycle configuration to transition data to Amazon S3 Glacier after 14 days
- B. Launch Amazon EC2 instances across two Availability Zones and place them behind an Elastic Load Balancer to ingest the alerts Create a script on the EC2 instances that will store the alerts in an Amazon S3 bucket Set up an S3 Lifecycle configuration to transition data to Amazon S3 Glacier after 14 days
- C. Create an Amazon Kinesis Data Firehose delivery stream to ingest the alerts Configure the Kinesis Data Firehose stream to deliver the alerts to an Amazon Elasticsearch Service (Amazon ES) cluster Set up the Amazon ES cluster to take manual snapshots every day and delete data from the cluster that is older than 14 days
- D. Create an Amazon Simple Queue Service (Amazon SQS) standard queue to ingest the alerts and set the message retention period to 14 days Configure consumers to poll the SQS queue check the age of the message and analyze the message data as needed If the message is 14 days old the consumer should

copy the message to an Amazon S3 bucket and delete the message from the SQS queue

Answer: A

Explanation:

<https://aws.amazon.com/kinesis/data-firehose/features/?nc=sn&loc=2#:~:text=into%20Amazon%20S3%2C%20Amazon%20Redshift%2C%20Amazon%20OpenSearch%20Service%2C%20Kinesis,Delivery%20streams>

NEW QUESTION 89

- (Topic 1)

A company is designing an application. The application uses an AWS Lambda function to receive information through Amazon API Gateway and to store the information in an Amazon Aurora PostgreSQL database.

During the proof-of-concept stage, the company has to increase the Lambda quotas significantly to handle the high volumes of data that the company needs to load into the database. A solutions architect must recommend a new design to improve scalability and minimize the configuration effort.

Which solution will meet these requirements?

- A. Refactor the Lambda function code to Apache Tomcat code that runs on Amazon EC2 instance
- B. Connect the database by using native Java Database Connectivity (JDBC) drivers.
- C. Change the platform from Aurora to Amazon DynamoDB
- D. Provision a DynamoDB Accelerator (DAX) cluster
- E. Use the DAX client SDK to point the existing DynamoDB API calls at the DAX cluster.
- F. Set up two Lambda function
- G. Configure one function to receive the information
- H. Configure the other function to load the information into the database
- I. Integrate the Lambda functions by using Amazon Simple Notification Service (Amazon SNS).
- J. Set up two Lambda function
- K. Configure one function to receive the information
- L. Configure the other function to load the information into the database
- M. Integrate the Lambda functions by using an Amazon Simple Queue Service (Amazon SQS) queue.

Answer: B

Explanation:

bottlenecks can be avoided with queues (SQS).

NEW QUESTION 90

- (Topic 1)

A company uses AWS Organizations to manage multiple AWS accounts for different departments. The management account has an Amazon S3 bucket that contains project reports. The company wants to limit access to this S3 bucket to only users of accounts within the organization in AWS Organizations.

Which solution meets these requirements with the LEAST amount of operational overhead?

- A. Add the aws:PrincipalOrgID global condition key with a reference to the organization ID to the S3 bucket policy.
- B. Create an organizational unit (OU) for each department
- C. Add the aws:PrincipalOrgPaths global condition key to the S3 bucket policy.
- D. Use AWS CloudTrail to monitor the CreateAccount, InviteAccountToOrganization, LeaveOrganization, and RemoveAccountFromOrganization event
- E. Update the S3 bucket policy accordingly.
- F. Tag each user that needs access to the S3 bucket
- G. Add the aws:PrincipalTag global condition key to the S3 bucket policy.

Answer: A

Explanation:

<https://aws.amazon.com/blogs/security/control-access-to-aws-resources-by-using-the-aws-organization-of-iam-principals/>

The aws:PrincipalOrgID global key provides an alternative to listing all the account IDs for all AWS accounts in an organization. For example, the following Amazon S3 bucket policy allows members of any account in the XXX organization to add an object into the examtopics bucket.

```
{
  "Version": "2020-09-10",
  "Statement": [
    {
      "Sid": "AllowPutObject",
      "Effect": "Allow",
      "Principal": "*",
      "Action": "s3:PutObject",
      "Resource": "arn:aws:s3:::examtopics/*",
      "Condition": {
        "StringEquals": {
          "aws:PrincipalOrgID": ["XXX"]
        }
      }
    }
  ]
}
```

https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_condition-keys.html

NEW QUESTION 95

- (Topic 1)

A company runs its Infrastructure on AWS and has a registered base of 700,000 users for a document management application. The company intends to create a product that converts large PDF files to JPG image files. The PDF files average 5 MB in size. The company needs to store the original files and the converted files. A solutions architect must design a scalable solution to accommodate demand that will grow rapidly over time.

Which solution meets these requirements MOST cost-effectively?

- A. Save the PDF files to Amazon S3. Configure an S3 PUT event to invoke an AWS Lambda function to convert the files to JPG format and store them back in Amazon S3.
- B. Save the PDF files to Amazon DynamoDB.
- C. Use the DynamoDB Streams feature to invoke an AWS Lambda function to convert the files to JPG format and store them back in DynamoDB.
- D. Upload the PDF files to an AWS Elastic Beanstalk application that includes Amazon EC2 instances.
- E. Amazon Elastic Block Store (Amazon EBS) storage and an Auto Scaling group.
- F. Use a program in the EC2 instances to convert the files to JPG format. Save the PDF files and the JPG files in the EBS store.
- G. Upload the PDF files to an AWS Elastic Beanstalk application that includes Amazon EC2 instances, Amazon Elastic File System (Amazon EFS) storage, and an

Auto Scaling group

H. Use a program in the EC2 instances to convert the file to jpg format Save the pdf files and the jpg files in the EBS store.

Answer: A

Explanation:

Elastic BeanStalk is expensive, and DocumentDB has a 400KB max to upload files. So Lambda and S3 should be the one.

NEW QUESTION 99

- (Topic 1)

A company wants to reduce the cost of its existing three-tier web architecture. The web, application, and database servers are running on Amazon EC2 instances for the development, test, and production environments. The EC2 instances average 30% CPU utilization during peak hours and 10% CPU utilization during non-peak hours.

The production EC2 instances run 24 hours a day. The development and test EC2 instances run for at least 8 hours each day. The company plans to implement automation to stop the development and test EC2 instances when they are not in use.

Which EC2 instance purchasing solution will meet the company's requirements MOST cost-effectively?

- A. Use Spot Instances for the production EC2 instance
- B. Use Reserved Instances for the development and test EC2 instances.
- C. Use Reserved Instances for the production EC2 instance
- D. Use On-Demand Instances for the development and test EC2 instances.
- E. Use Spot blocks for the production EC2 instance
- F. Use Reserved Instances for the development and test EC2 instances.
- G. Use On-Demand Instances for the production EC2 instance
- H. Use Spot blocks for the development and test EC2 instances.

Answer: B

NEW QUESTION 101

- (Topic 2)

A company hosts a two-tier application on Amazon EC2 instances and Amazon RDS. The application's demand varies based on the time of day. The load is minimal after work hours and on weekends. The EC2 instances run in an EC2 Auto Scaling group that is configured with a minimum of two instances and a maximum of five instances. The application must be available at all times, but the company is concerned about overall cost.

Which solution meets the availability requirement MOST cost-effectively?

- A. Use all EC2 Spot Instance
- B. Stop the RDS database when it is not in use.
- C. Purchase EC2 Instance Savings Plans to cover five EC2 instance
- D. Purchase an RDS Reserved DB Instance
- E. Purchase two EC2 Reserved Instances Use up to three additional EC2 Spot Instances as needed
- F. Stop the RDS database when it is not in use.
- G. Purchase EC2 Instance Savings Plans to cover two EC2 instance
- H. Use up to three additional EC2 On-Demand Instances as needed
- I. Purchase an RDS Reserved DB Instance.

Answer: C

Explanation:

This solution meets the requirements of a two-tier application that has a variable demand based on the time of day and must be available at all times, while minimizing the overall cost. EC2 Reserved Instances can provide significant savings compared to On-Demand Instances for the baseline level of usage, and they can guarantee capacity reservation when needed. EC2 Spot Instances can provide up to 90% savings compared to On-Demand Instances for any additional capacity that the application needs during peak hours. Spot Instances are suitable for stateless applications that can tolerate interruptions and can be replaced by other instances. Stopping the RDS database when it is not in use can reduce the cost of running the database tier.

Option A is incorrect because using all EC2 Spot Instances can affect the availability of the application if there are not enough spare capacity or if the Spot price exceeds the maximum price. Stopping the RDS database when it is not in use can reduce the cost of running the database tier, but it can also affect the availability of the application. Option B is incorrect because purchasing EC2 Instance Savings Plans to cover five EC2 instances can lock in a fixed amount of compute usage per hour, which may not match the actual usage pattern of the application. Purchasing an RDS Reserved DB Instance can provide savings for the database tier, but it does not allow stopping the database when it is not in use. Option D is incorrect because purchasing EC2 Instance Savings Plans to cover two EC2 instances can lock in a fixed amount of compute usage per hour, which may not match the actual usage pattern of the application. Using up to three additional EC2 On-Demand Instances as needed can incur higher costs than using Spot Instances.

References:

? <https://aws.amazon.com/ec2/pricing/reserved-instances/>

? <https://aws.amazon.com/ec2/spot/>

? https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_StopInstance.html

NEW QUESTION 103

- (Topic 2)

A company wants to migrate its on-premises data center to AWS. According to the company's compliance requirements, the company can use only the ap-northeast-3 Region. Company administrators are not permitted to connect VPCs to the internet.

Which solutions will meet these requirements? (Choose two.)

- A. Use AWS Control Tower to implement data residency guardrails to deny internet access and deny access to all AWS Regions except ap-northeast-3.
- B. Use rules in AWS WAF to prevent internet access
- C. Deny access to all AWS Regions except ap-northeast-3 in the AWS account settings.
- D. Use AWS Organizations to configure service control policies (SCPs) that prevent VPCs from gaining internet access
- E. Deny access to all AWS Regions except ap-northeast-3.
- F. Create an outbound rule for the network ACL in each VPC to deny all traffic from 0.0.0.0/0. Create an IAM policy for each user to prevent the use of any AWS Region other than ap-northeast-3.
- G. Use AWS Config to activate managed rules to detect and alert for internet gateways and to detect and alert for new resources deployed outside of ap-northeast-3.

Answer: AC

Explanation:

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_vpc.html#example_vpc_2

NEW QUESTION 105

- (Topic 2)

A company has a multi-tier application that runs six front-end web servers in an Amazon EC2 Auto Scaling group in a single Availability Zone behind an Application Load Balancer (ALB). A solutions architect needs to modify the infrastructure to be highly available without modifying the application. Which architecture should the solutions architect choose that provides high availability?

- A. Create an Auto Scaling group that uses three instances across each of two Regions.
- B. Modify the Auto Scaling group to use three instances across each of two Availability Zones.
- C. Create an Auto Scaling template that can be used to quickly create more instances in another Region.
- D. Change the ALB in front of the Amazon EC2 instances in a round-robin configuration to balance traffic to the web tier.

Answer: B

Explanation:

High availability can be enabled for this architecture quite simply by modifying the existing Auto Scaling group to use multiple availability zones. The ASG will automatically balance the load so you don't actually need to specify the instances per AZ.

NEW QUESTION 110

- (Topic 2)

A company wants to use the AWS Cloud to make an existing application highly available and resilient. The current version of the application resides in the company's data center. The application recently experienced data loss after a database server crashed because of an unexpected power outage. The company needs a solution that avoids any single points of failure. The solution must give the application the ability to scale to meet user demand. Which solution will meet these requirements?

- A. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones.
- B. Use an Amazon RDS DB instance in a Multi-AZ configuration.
- C. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group in a single Availability Zone.
- D. Deploy the database on an EC2 instance.
- E. Enable EC2 Auto Recovery.
- F. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones.
- G. Use an Amazon RDS DB instance with a read replica in a single Availability Zone.
- H. Promote the read replica to replace the primary DB instance if the primary DB instance fails.
- I. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones. Deploy the primary and secondary database servers on EC2 instances across multiple Availability Zones. Use Amazon Elastic Block Store (Amazon EBS) Multi-Attach to create shared storage between the instances.

Answer: A

Explanation:

Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones. Use an Amazon RDS DB instance in a Multi-AZ configuration. To make an existing application highly available and resilient while avoiding any single points of failure and giving the application the ability to scale to meet user demand, the best solution would be to deploy the application servers using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones and use an Amazon RDS DB instance in a Multi-AZ configuration. By using an Amazon RDS DB instance in a Multi-AZ configuration, the database is automatically replicated across multiple Availability Zones, ensuring that the database is highly available and can withstand the failure of a single Availability Zone. This provides fault tolerance and avoids any single points of failure.

NEW QUESTION 112

- (Topic 2)

An ecommerce company hosts its analytics application in the AWS Cloud. The application generates about 300 MB of data each month. The data is stored in JSON format. The company is evaluating a disaster recovery solution to back up the data. The data must be accessible in milliseconds if it is needed, and the data must be kept for 30 days.

Which solution meets these requirements MOST cost-effectively?

- A. Amazon OpenSearch Service (Amazon Elasticsearch Service)
- B. Amazon S3 Glacier
- C. Amazon S3 Standard
- D. Amazon RDS for PostgreSQL

Answer: C

Explanation:

This solution meets the requirements of a disaster recovery solution to back up the data that is generated by an analytics application, stored in JSON format, and must be accessible in milliseconds if it is needed. Amazon S3 Standard is a durable and scalable storage class for frequently accessed data. It can store any amount of data and provide high availability and performance. It can also support millisecond access time for data retrieval.

Option A is incorrect because Amazon OpenSearch Service (Amazon Elasticsearch Service) is a search and analytics service that can index and query data, but it is not a backup solution for data stored in JSON format. Option B is incorrect because Amazon S3 Glacier is a low-cost storage class for data archiving and long-term backup, but it does not support millisecond access time for data retrieval. Option D is incorrect because Amazon RDS for PostgreSQL is a relational database service that can store and query structured data, but it is not a backup solution for data stored in JSON format.

References:

? <https://aws.amazon.com/s3/storage-classes/>

? https://aws.amazon.com/s3/faqs/#Durability_and_data_protection

NEW QUESTION 116

- (Topic 2)

A company is migrating its on-premises PostgreSQL database to Amazon Aurora PostgreSQL. The on-premises database must remain online and accessible during the migration. The Aurora database must remain synchronized with the on-premises database. Which combination of actions must a solutions architect take to meet these requirements? (Choose two.)

- A. Create an ongoing replication task.
- B. Create a database backup of the on-premises database
- C. Create an AWS Database Migration Service (AWS DMS) replication server
- D. Convert the database schema by using the AWS Schema Conversion Tool (AWS SCT).
- E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to monitor the database synchronization

Answer: AC

Explanation:

AWS Database Migration Service supports homogeneous migrations such as Oracle to Oracle, as well as heterogeneous migrations between different database platforms, such as Oracle or Microsoft SQL Server to Amazon Aurora. With AWS Database Migration Service, you can also continuously replicate data with low latency from any supported source to any supported target. For example, you can replicate from multiple sources to Amazon Simple Storage Service (Amazon S3) to build a highly available and scalable data lake solution. You can also consolidate databases into a petabyte-scale data warehouse by streaming data to Amazon Redshift. Learn more about the supported source and target databases. <https://aws.amazon.com/dms/>

NEW QUESTION 117

- (Topic 2)

A company is building a web-based application running on Amazon EC2 instances in multiple Availability Zones. The web application will provide access to a repository of text documents totaling about 900 TB in size. The company anticipates that the web application will experience periods of high demand. A solutions architect must ensure that the storage component for the text documents can scale to meet the demand of the application at all times. The company is concerned about the overall cost of the solution.

Which storage solution meets these requirements MOST cost-effectively?

- A. Amazon Elastic Block Store (Amazon EBS)
- B. Amazon Elastic File System (Amazon EFS)
- C. Amazon Elasticsearch Service (Amazon ES)
- D. Amazon S3

Answer: D

Explanation:

Amazon S3 is cheapest and can be accessed from anywhere.

NEW QUESTION 120

- (Topic 2)

An online retail company has more than 50 million active customers and receives more than 25,000 orders each day. The company collects purchase data for customers and stores this data in Amazon S3. Additional customer data is stored in Amazon RDS.

The company wants to make all the data available to various teams so that the teams can perform analytics. The solution must provide the ability to manage fine-grained permissions for the data and must minimize operational overhead.

Which solution will meet these requirements?

- A. Migrate the purchase data to write directly to Amazon RD
- B. Use RDS access controls to limit access.
- C. Schedule an AWS Lambda function to periodically copy data from Amazon RDS to Amazon S3. Create an AWS Glue crawle
- D. Use Amazon Athena to query the dat
- E. Use S3 policies to limit access.
- F. Create a data lake by using AWS Lake Formatio
- G. Create an AWS Glue JDBC connection to Amazon RD
- H. Register (he S3 bucket in Lake Formatio
- I. Use Lake Formation access controls to limit access.
- J. Create an Amazon Redshift cluste
- K. Schedule an AWS Lambda function to periodically copy data from Amazon S3 and Amazon RDS to Amazon Redshif
- L. Use Amazon Redshift access controls to limit access.

Answer: C

Explanation:

To make all the data available to various teams and minimize operational overhead, the company can create a data lake by using AWS Lake Formation. This will allow the company to centralize all the data in one place and use fine-grained access controls to manage access to the data. To meet the requirements of the company, the solutions architect can create a data lake by using AWS Lake Formation, create an AWS Glue JDBC connection to Amazon RDS, and register the S3 bucket in Lake Formation. The solutions architect can then use Lake Formation access controls to limit access to the data. This solution will provide the ability to manage fine-grained permissions for the data and minimize operational overhead.

NEW QUESTION 122

- (Topic 2)

A company is concerned about the security of its public web application due to recent web attacks. The application uses an Application Load Balancer (ALB). A solutions architect must reduce the risk of DDoS attacks against the application.

What should the solutions architect do to meet this requirement?

- A. Add an Amazon Inspector agent to the ALB.
- B. Configure Amazon Macie to prevent attacks.
- C. Enable AWS Shield Advanced to prevent attacks.
- D. Configure Amazon GuardDuty to monitor the ALB.

Answer: C

Explanation:

AWS Shield Advanced provides expanded DDoS attack protection for your Amazon EC2 instances, Elastic Load Balancing load balancers, CloudFront distributions, Route 53 hosted zones, and AWS Global Accelerator standard accelerators. <https://docs.aws.amazon.com/waf/latest/developerguide/what-is-aws-waf.html>

NEW QUESTION 125

- (Topic 2)

A new employee has joined a company as a deployment engineer. The deployment engineer will be using AWS CloudFormation templates to create multiple AWS resources. A solutions architect wants the deployment engineer to perform job activities while following the principle of least privilege. Which steps should the solutions architect do in conjunction to reach this goal? (Select two.)

- A. Have the deployment engineer use AWS account root user credentials for performing AWS CloudFormation stack operations.
- B. Create a new IAM user for the deployment engineer and add the IAM user to a group that has the PowerUsers IAM policy attached.
- C. Create a new IAM user for the deployment engineer and add the IAM user to a group that has the Administrate/Access IAM policy attached.
- D. Create a new IAM User for the deployment engineer and add the IAM user to a group that has an IAM policy that allows AWS CloudFormation actions only.
- E. Create an IAM role for the deployment engineer to explicitly define the permissions specific to the AWS CloudFormation stack and launch stacks using Dial IAM role.

Answer: DE

Explanation:

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html https://docs.aws.amazon.com/IAM/latest/UserGuide/id_users.html

NEW QUESTION 127

- (Topic 2)

A company is planning to build a high performance computing (HPC) workload as a service solution that is hosted on AWS. A group of 16 Amazon EC2 Linux instances requires the lowest possible latency for node-to-node communication. The instances also need a shared block device volume for high-performing storage.

Which solution will meet these requirements?

- A. Use a duster placement group
- B. Attach a single Provisioned IOPS SSD Amazon Elastic Block Store (Amazon EBS) volume to all the instances by using Amazon EBS Multi-Attach
- C. Use a cluster placement group
- D. Create shared file systems across the instances by using Amazon Elastic File System (Amazon EFS)
- E. Use a partition placement group
- F. Create shared tile systems across the instances by using Amazon Elastic File System (Amazon EFS).
- G. Use a spread placement group
- H. Attach a single Provisioned IOPS SSD Amazon Elastic Block Store (Amazon EBS) volume to all the instances by using Amazon EBS Multi-Attach

Answer: A

Explanation:

1. lowest possible latency + node to node ==> cluster placement(must be within one AZ), so C, D out
 - * 2. For EBS Multi-Attach, up to 16 instances can be attached to a single volume==>we have 16 linux instance==>more close to A
 - * 3. "need a shared block device volume"==>EBS Multi-attach is Block Storage whereas EFS is File Storage==> B out
 - * 4. EFS automatically replicates data within and across 3 AZ==>we use cluster placement
- so all EC2 are within one AZ.
- * 5. EBS Multi-attach volumes can be used for clients within a single AZ. <https://repost.aws/questions/QUK2RANw1QTKCwpDUwCCI72A/efs-vs-efs-mult-attach>

NEW QUESTION 131

- (Topic 2)

A company wants to move its application to a serverless solution. The serverless solution needs to analyze existing and new data by using SL. The company stores the data in an Amazon S3 bucket. The data requires encryption and must be replicated to a different AWS Region. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create a new S3 bucket
- B. Load the data into the new S3 bucket
- C. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Region
- D. Use server-side encryption with AWS KMS multi-Region keys (SSE-KMS). Use Amazon Athena to query the data.
- E. Create a new S3 bucket
- F. Load the data into the new S3 bucket
- G. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Region
- H. Use server-side encryption with AWS KMS multi-Region keys (SSE-KMS). Use Amazon RDS to query the data.
- I. Load the data into the existing S3 bucket
- J. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Region
- K. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Use Amazon Athena to query the data.
- L. Load the data into the existing S3 bucket
- M. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Region
- N. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Use Amazon RDS to query the data.

Answer: A

Explanation:

This solution meets the requirements of a serverless solution, encryption, replication, and SQL analysis with the least operational overhead. Amazon Athena is a serverless interactive query service that can analyze data in S3 using standard SQL. S3 Cross-Region Replication (CRR) can replicate encrypted objects to an S3 bucket in another Region automatically. Server-side encryption with AWS KMS multi-Region keys (SSE-KMS) can encrypt the data at rest using keys that are replicated across multiple Regions. Creating a new S3 bucket can avoid potential conflicts with existing data or configurations. Option B is incorrect because Amazon RDS is not a serverless solution and it cannot query data in S3 directly. Option C is incorrect because server-side encryption with Amazon S3 managed encryption keys (SSE-S3) does not use KMS keys and it does not support multi-Region replication. Option D is incorrect because Amazon RDS is not a serverless solution and it cannot query data in S3 directly. It is also incorrect for the same reason as option C. References:

- ? <https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-walkthrough-4.html>
- ? <https://aws.amazon.com/blogs/storage/considering-four-different-replication-options-for-data-in-amazon-s3/>
- ? <https://docs.aws.amazon.com/AmazonS3/latest/userguide/UsingEncryption.html>
- ? <https://aws.amazon.com/athena/>

NEW QUESTION 135

- (Topic 2)

A company has an AWS account used for software engineering. The AWS account has access to the company's on-premises data center through a pair of AWS Direct Connect connections. All non-VPC traffic routes to the virtual private gateway.

A development team recently created an AWS Lambda function through the console. The development team needs to allow the function to access a database that runs in a private subnet in the company's data center.

Which solution will meet these requirements?

- A. Configure the Lambda function to run in the VPC with the appropriate security group.
- B. Set up a VPN connection from AWS to the data center.
- C. Route the traffic from the Lambda function through the VPN.
- D. Update the route tables in the VPC to allow the Lambda function to access the on-premises data center through Direct Connect.
- E. Create an Elastic IP address.
- F. Configure the Lambda function to send traffic through the Elastic IP address without an elastic network interface.

Answer: A

Explanation:

<https://docs.aws.amazon.com/lambda/latest/dg/configuration-vpc.html#vpc-managing-eni>

NEW QUESTION 140

- (Topic 2)

A company runs a production application on a fleet of Amazon EC2 instances. The application reads the data from an Amazon SQS queue and processes the messages in parallel. The message volume is unpredictable and often has intermittent traffic. This application should continually process messages without any downtime.

Which solution meets these requirements MOST cost-effectively?

- A. Use Spot Instances exclusively to handle the maximum capacity required.
- B. Use Reserved Instances exclusively to handle the maximum capacity required.
- C. Use Reserved Instances for the baseline capacity and use Spot Instances to handle additional capacity.
- D. Use Reserved Instances for the baseline capacity and use On-Demand Instances to handle additional capacity.

Answer: D

Explanation:

We recommend that you use On-Demand Instances for applications with short-term, irregular workloads that cannot be interrupted.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-on-demand-instances.html>

NEW QUESTION 142

- (Topic 2)

A gaming company hosts a browser-based application on AWS. The users of the application consume a large number of videos and images that are stored in Amazon S3. This content is the same for all users.

The application has increased in popularity, and millions of users worldwide are accessing these media files. The company wants to provide the files to the users while reducing the load on the origin.

Which solution meets these requirements MOST cost-effectively?

- A. Deploy an AWS Global Accelerator accelerator in front of the web servers.
- B. Deploy an Amazon CloudFront web distribution in front of the S3 bucket.
- C. Deploy an Amazon ElastiCache for Redis instance in front of the web servers.
- D. Deploy an Amazon ElastiCache for Memcached instance in front of the web servers.

Answer: B

Explanation:

ElastiCache, enhances the performance of web applications by quickly retrieving information from fully-managed in-memory data stores. It utilizes Memcached and Redis, and manages to considerably reduce the time your applications would, otherwise, take to read data from disk-based databases. Amazon CloudFront supports dynamic content from HTTP and WebSocket protocols, which are based on the Transmission Control Protocol (TCP) protocol. Common use cases include dynamic API calls, web pages and web applications, as well as an application's static files such as audio and images. It also supports on-demand media streaming over HTTP. AWS Global Accelerator supports both User Datagram Protocol (UDP) and TCP-based protocols. It is commonly used for non-HTTP use cases, such as gaming, IoT and voice over IP. It is also good for HTTP use cases that need static IP addresses or fast regional failover.

NEW QUESTION 143

- (Topic 2)

A company has a highly dynamic batch processing job that uses many Amazon EC2 instances to complete it. The job is stateless in nature, can be started and stopped at any given time with no negative impact, and typically takes upwards of 60 minutes total to complete. The company has asked a solutions architect to design a scalable and cost-effective solution that meets the requirements of the job.

What should the solutions architect recommend?

- A. Implement EC2 Spot Instances
- B. Purchase EC2 Reserved Instances
- C. Implement EC2 On-Demand Instances
- D. Implement the processing on AWS Lambda

Answer: A

Explanation:

EC2 Spot Instances allow users to bid on spare Amazon EC2 computing capacity and can be a cost-effective solution for stateless, interruptible workloads that can be started and stopped at any time. Since the batch processing job is stateless, can be started and stopped at any time, and typically takes upwards of 60 minutes to complete, EC2 Spot Instances would be a good fit for this workload.

NEW QUESTION 146

- (Topic 2)

A company runs an Oracle database on premises. As part of the company's migration to AWS, the company wants to upgrade the database to the most recent available version. The company also wants to set up disaster recovery (DR) for the database. The company needs to minimize the operational overhead for normal operations and DR setup. The company also needs to maintain access to the database's underlying operating system. Which solution will meet these requirements?

- A. Migrate the Oracle database to an Amazon EC2 instance
- B. Set up database replication to a different AWS Region.
- C. Migrate the Oracle database to Amazon RDS for Oracle
- D. Activate Cross-Region automated backups to replicate the snapshots to another AWS Region.
- E. Migrate the Oracle database to Amazon RDS Custom for Oracle
- F. Create a read replica for the database in another AWS Region.
- G. Migrate the Oracle database to Amazon RDS for Oracle
- H. Create a standby database in another Availability Zone.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/rds-custom.html> and <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/working-with-custom-oracle.html>

NEW QUESTION 147

- (Topic 2)

An ecommerce company has an order-processing application that uses Amazon API Gateway and an AWS Lambda function. The application stores data in an Amazon Aurora PostgreSQL database. During a recent sales event, a sudden surge in customer orders occurred. Some customers experienced timeouts and the application did not process the orders of those customers. A solutions architect determined that the CPU utilization and memory utilization were high on the database because of a large number of open connections. The solutions architect needs to prevent the timeout errors while making the least possible changes to the application.

Which solution will meet these requirements?

- A. Configure provisioned concurrency for the Lambda function. Modify the database to be a global database in multiple AWS Regions.
- B. Use Amazon RDS Proxy to create a proxy for the database. Modify the Lambda function to use the RDS Proxy endpoint instead of the database endpoint.
- C. Create a read replica for the database in a different AWS Region. Use query string parameters in API Gateway to route traffic to the read replica.
- D. Migrate the data from Aurora PostgreSQL to Amazon DynamoDB by using AWS Database Migration Service (AWS DMS). Modify the Lambda function to use the DynamoDB table.

Answer: B

Explanation:

Many applications, including those built on modern serverless architectures, can have a large number of open connections to the database server and may open and close database connections at a high rate, exhausting database memory and compute resources. Amazon RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability. <https://aws.amazon.com/id/rds/proxy/>

NEW QUESTION 151

- (Topic 2)

A solutions architect is optimizing a website for an upcoming musical event. Videos of the performances will be streamed in real time and then will be available on demand. The event is expected to attract a global online audience.

Which service will improve the performance of both the real-time and on-demand streaming?

- A. Amazon CloudFront
- B. AWS Global Accelerator
- C. Amazon Route 53
- D. Amazon S3 Transfer Acceleration

Answer: A

Explanation:

You can use CloudFront to deliver video on demand (VOD) or live streaming video using any HTTP origin. One way you can set up video workflows in the cloud is by using CloudFront together with AWS Media Services. <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/on-demand-streaming-video.html>

NEW QUESTION 154

- (Topic 2)

A company wants to run a gaming application on Amazon EC2 instances that are part of an Auto Scaling group in the AWS Cloud. The application will transmit data by using UDP packets. The company wants to ensure that the application can scale out and in as traffic increases and decreases.

What should a solutions architect do to meet these requirements?

- A. Attach a Network Load Balancer to the Auto Scaling group.
- B. Attach an Application Load Balancer to the Auto Scaling group.
- C. Deploy an Amazon Route 53 record set with a weighted policy to route traffic appropriately.
- D. Deploy a NAT instance that is configured with port forwarding to the EC2 instances in the Auto Scaling group.

Answer: A

Explanation:

This solution meets the requirements of running a gaming application that transmits data by using UDP packets and scaling out and in as traffic increases and decreases. A Network Load Balancer can handle millions of requests per second while maintaining high throughput at ultra low latency, and it supports both TCP and UDP protocols. An Auto Scaling group can automatically adjust the number of EC2 instances based on the demand and the scaling policies. Option B is incorrect because an Application Load Balancer does not support UDP protocol, only HTTP and HTTPS. Option C is incorrect because Amazon Route 53 is a DNS service that can route traffic based on different policies, but it does not provide load balancing or scaling capabilities. Option D is incorrect because a NAT instance is used to enable instances in a private subnet to connect to the internet or other AWS services, but it does not provide load balancing or scaling capabilities.

References:

- ? <https://aws.amazon.com/blogs/aws/new-udp-load-balancing-for-network-load-balancer/>
- ? <https://docs.aws.amazon.com/autoscaling/ec2/userguide/AutoScalingGroup.html>

NEW QUESTION 159

- (Topic 2)

A company has a data ingestion workflow that includes the following components:

- An Amazon Simple Notification Service (Amazon SNS) topic that receives notifications about new data deliveries
- An AWS Lambda function that processes and stores the data

The ingestion workflow occasionally fails because of network connectivity issues. When a failure occurs the corresponding data is not ingested unless the company manually reruns the job. What should a solutions architect do to ensure that all notifications are eventually processed?

- A. Configure the Lambda function (or deployment across multiple Availability Zones)
- B. Modify the Lambda functions configuration to increase the CPU and memory allocations for the function
- C. Configure the SNS topic's retry strategy to increase both the number of retries and the wait time between retries
- D. Configure an Amazon Simple Queue Service (Amazon SQS) queue as the on-failure destination. Modify the Lambda function to process messages in the queue

Answer: D

Explanation:

<https://docs.aws.amazon.com/sns/latest/dg/sns-dead-letter-queues.html>

NEW QUESTION 162

- (Topic 2)

A company has two applications: a sender application that sends messages with payloads to be processed and a processing application intended to receive the messages with payloads. The company wants to implement an AWS service to handle messages between the two applications. The sender application can send about 1,000 messages each hour. The messages may take up to 2 days to be processed. If the messages fail to process, they must be retained so that they do not impact the processing of any remaining messages.

Which solution meets these requirements and is the MOST operationally efficient?

- A. Set up an Amazon EC2 instance running a Redis database
- B. Configure both applications to use the instance
- C. Store, process, and delete the messages, respectively.
- D. Use an Amazon Kinesis data stream to receive the messages from the sender application
- E. Integrate the processing application with the Kinesis Client Library (KCL).
- F. Integrate the sender and processor applications with an Amazon Simple Queue Service (Amazon SQS) queue
- G. Configure a dead-letter queue to collect the messages that failed to process.
- H. Subscribe the processing application to an Amazon Simple Notification Service (Amazon SNS) topic to receive notifications to process
- I. Integrate the sender application to write to the SNS topic.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/compute/building-loosely-coupled-scalable-c-applications-with-amazon-sqs-and-amazon-sns/>
<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-dead-letter-queues.html>

NEW QUESTION 167

- (Topic 2)

A company has a service that produces event data. The company wants to use AWS to process the event data as it is received. The data is written in a specific order that must be maintained throughout processing. The company wants to implement a solution that minimizes operational overhead.

How should a solutions architect accomplish this?

- A. Create an Amazon Simple Queue Service (Amazon SQS) FIFO queue to hold messages. Set up an AWS Lambda function to process messages from the queue.
- B. Create an Amazon Simple Notification Service (Amazon SNS) topic to deliver notifications containing payloads to process. Configure an AWS Lambda function as a subscriber.
- C. Create an Amazon Simple Queue Service (Amazon SQS) standard queue to hold messages.
- D. Set up an AWS Lambda function to process messages from the queue independently.
- E. Create an Amazon Simple Notification Service (Amazon SNS) topic to deliver notifications containing payloads to process.
- F. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a subscriber.

Answer: A

Explanation:

The details are revealed in the below URL: <https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/FIFO-queues.html>

FIFO (First-In-First-Out) queues are designed to enhance messaging between applications when the order of operations and events is critical, or where duplicates can't be tolerated. Examples of situations where you might use FIFO queues include the following: To make sure that user-entered commands are run in the right order. To display the correct product price by sending price modifications in the right order. To prevent a student from enrolling in a course before registering for an account.

NEW QUESTION 169

- (Topic 2)

A company is developing a file-sharing application that will use an Amazon S3 bucket for storage. The company wants to serve all the files through an Amazon CloudFront distribution. The company does not want the files to be accessible through direct navigation to the S3 URL. What should a solutions architect do to meet these requirements?

- A. Write individual policies for each S3 bucket to grant read permission for only CloudFront access.
- B. Create an IAM user
- C. Grant the user read permission to objects in the S3 bucket
- D. Assign the user to CloudFront.
- E. Write an S3 bucket policy that assigns the CloudFront distribution ID as the Principal and assigns the target S3 bucket as the Amazon Resource Name (ARN).
- F. Create an origin access identity (OAI). Assign the OAI to the CloudFront distribution
- G. Configure the S3 bucket permissions so that only the OAI has read permission.

Answer: D

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-access-to-amazon-s3/>

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3.html#private-content-restricting-access-to-s3-overview>

NEW QUESTION 173

- (Topic 2)

A company runs a stateless web application in production on a group of Amazon EC2 On-Demand Instances behind an Application Load Balancer. The application experiences heavy usage during an 8-hour period each business day. Application usage is moderate and steady overnight. Application usage is low during weekends.

The company wants to minimize its EC2 costs without affecting the availability of the application.

Which solution will meet these requirements?

- A. Use Spot Instances for the entire workload.
- B. Use Reserved instances for the baseline level of usage. Use Spot Instances for any additional capacity that the application needs.
- C. Use On-Demand Instances for the baseline level of usage.
- D. Use Spot Instances for any additional capacity that the application needs.
- E. Use Dedicated Instances for the baseline level of usage.
- F. Use On-Demand Instances for any additional capacity that the application needs.

Answer: B

Explanation:

Reserved is cheaper than on demand the company has. And it meets the availability (HA) requirement as to spot instance that can be disrupted at any time. PRICING BELOW. On-Demand: 0% There's no commitment from you. You pay the most with this option. Reserved : 40%-60% 1-year or 3-year commitment from you. You save money from that commitment. Spot 50%-90% Ridiculously inexpensive because there's no commitment from the AWS side.

NEW QUESTION 178

- (Topic 2)

A company needs to retain application logs files for a critical application for 10 years. The application team regularly accesses logs from the past month for troubleshooting, but logs older than 1 month are rarely accessed. The application generates more than 10 TB of logs per month.

Which storage option meets these requirements MOST cost-effectively?

- A. Store the logs in Amazon S3. Use AWS Backup to move logs more than 1 month old to S3 Glacier Deep Archive.
- B. Store the logs in Amazon S3. Use S3 Lifecycle policies to move logs more than 1 month old to S3 Glacier Deep Archive.
- C. Store the logs in Amazon CloudWatch Logs. Use AWS Backup to move logs more than 1 month old to S3 Glacier Deep Archive.
- D. Store the logs in Amazon CloudWatch Logs. Use Amazon S3 Lifecycle policies to move logs more than 1 month old to S3 Glacier Deep Archive.

Answer: B

Explanation:

You need S3 to be able to archive the logs after one month. Cannot do that with CloudWatch Logs.

NEW QUESTION 179

- (Topic 2)

A solutions architect needs to implement a solution to reduce a company's storage costs. All the company's data is in the Amazon S3 Standard storage class. The company must keep all data for at least 25 years. Data from the most recent 2 years must be highly available and immediately retrievable.

Which solution will meet these requirements?

- A. Set up an S3 Lifecycle policy to transition objects to S3 Glacier Deep Archive immediately.
- B. Set up an S3 Lifecycle policy to transition objects to S3 Glacier Deep Archive after 2 years.
- C. Use S3 Intelligent-Tiering.
- D. Activate the archiving option to ensure that data is archived in S3 Glacier Deep Archive.
- E. Set up an S3 Lifecycle policy to transition objects to S3 One Zone-Infrequent Access (S3 One Zone-IA) immediately and to S3 Glacier Deep Archive after 2 years.

Answer: B

Explanation:

https://aws.amazon.com/about-aws/whats-new/2018/04/announcing-s3-one-zone-infrequent-access-a-new-amazon-s3-storage-class/?nc1=h_ls

NEW QUESTION 182

- (Topic 2)

A solutions architect needs to securely store a database user name and password that an application uses to access an Amazon RDS DB instance. The application that accesses the database runs on an Amazon EC2 instance. The solutions architect wants to create a secure parameter in AWS Systems Manager

Parameter Store.

What should the solutions architect do to meet this requirement?

- A. Create an IAM role that has read access to the Parameter Store parameter
- B. Allow Decrypt access to an AWS Key Management Service (AWS KMS) key that is used to encrypt the parameter
- C. Assign this IAM role to the EC2 instance.
- D. Create an IAM policy that allows read access to the Parameter Store parameter
- E. Allow Decrypt access to an AWS Key Management Service (AWS KMS) key that is used to encrypt the parameter
- F. Assign this IAM policy to the EC2 instance.
- G. Create an IAM trust relationship between the Parameter Store parameter and the EC2 instance
- H. Specify Amazon RDS as a principal in the trust policy.
- I. Create an IAM trust relationship between the DB instance and the EC2 instance
- J. Specify Systems Manager as a principal in the trust policy.

Answer: B

Explanation:

https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_aws-services-that-work-with-iam.html

NEW QUESTION 186

- (Topic 2)

A company recently started using Amazon Aurora as the data store for its global ecommerce application. When large reports are run, developers report that the ecommerce application is performing poorly. After reviewing metrics in Amazon CloudWatch, a solutions architect finds that the ReadIOPS and CPU Utilization metrics are spiking when monthly reports run.

What is the MOST cost-effective solution?

- A. Migrate the monthly reporting to Amazon Redshift.
- B. Migrate the monthly reporting to an Aurora Replica
- C. Migrate the Aurora database to a larger instance class
- D. Increase the Provisioned IOPS on the Aurora instance

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Replication.html>

#Aurora.Replication.Replicas Aurora Replicas have two main purposes. You can issue queries to them to scale the read operations for your application. You typically do so by connecting to the reader endpoint of the cluster. That way, Aurora can spread the load for read-only connections across as many Aurora Replicas as you have in the cluster. Aurora Replicas also help to increase availability. If the writer instance in a cluster becomes unavailable, Aurora automatically promotes one of the reader instances to take its place as the new writer. <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Overview.html>

NEW QUESTION 189

- (Topic 2)

A company uses a three-tier web application to provide training to new employees. The application is accessed for only 12 hours every day. The company is using an Amazon RDS for MySQL DB instance to store information and wants to minimize costs.

What should a solutions architect do to meet these requirements?

- A. Configure an IAM policy for AWS Systems Manager Session Manager
- B. Create an IAM role for the policy
- C. Update the trust relationship of the role
- D. Set up automatic start and stop for the DB instance.
- E. Create an Amazon ElastiCache for Redis cache cluster that gives users the ability to access the data from the cache when the DB instance is stopped
- F. Invalidate the cache after the DB instance is started.
- G. Launch an Amazon EC2 instance
- H. Create an IAM role that grants access to Amazon RDS
- I. Attach the role to the EC2 instance
- J. Configure a cron job to start and stop the EC2 instance on the desired schedule.
- K. Create AWS Lambda functions to start and stop the DB instance
- L. Create Amazon EventBridge (Amazon CloudWatch Events) scheduled rules to invoke the Lambda function
- M. Configure the Lambda functions as event targets for the rules

Answer: D

Explanation:

In a typical development environment, dev and test databases are mostly utilized for 8 hours a day and sit idle when not in use. However, the databases are billed for the compute and storage costs during this idle time. To reduce the overall cost, Amazon RDS allows instances to be stopped temporarily. While the instance is stopped, you're charged for storage and backups, but not for the DB instance hours. Please note that a stopped instance will automatically be started after 7 days. This post presents a solution using AWS Lambda and Amazon EventBridge that allows you to schedule a Lambda function to stop and start the idle databases with specific tags to save on compute costs. The second post presents a solution that accomplishes stop and start of the idle Amazon RDS databases using AWS Systems Manager.

NEW QUESTION 190

- (Topic 2)

A company's website provides users with downloadable historical performance reports. The website needs a solution that will scale to meet the company's website demands globally. The solution should be cost-effective, limit the provisioning of infrastructure resources, and provide the fastest possible response time. Which combination should a solutions architect recommend to meet these requirements?

- A. Amazon CloudFront and Amazon S3
- B. AWS Lambda and Amazon DynamoDB
- C. Application Load Balancer with Amazon EC2 Auto Scaling
- D. Amazon Route 53 with internal Application Load Balancers

Answer: A

Explanation:

Cloudfront for rapid response and s3 to minimize infrastructure.

NEW QUESTION 195

- (Topic 2)

A company has implemented a self-managed DNS solution on three Amazon EC2 instances behind a Network Load Balancer (NLB) in the us-west-2 Region. Most of the company's users are located in the United States and Europe. The company wants to improve the performance and availability of the solution. The company launches and configures three EC2 instances in the eu-west-1 Region and adds the EC2 instances as targets for a new NLB.

Which solution can the company use to route traffic to all the EC2 instances?

- A. Create an Amazon Route 53 geolocation routing policy to route requests to one of the two NLB
- B. Create an Amazon CloudFront distributio
- C. Use the Route 53 record as the distribution's origin.
- D. Create a standard accelerator in AWS Global Accelerato
- E. Create endpoint groups in us- west-2 and eu-west-1. Add the two NLBs as endpoints for the endpoint groups.
- F. Attach Elastic IP addresses to the six EC2 instance
- G. Create an Amazon Route 53 geolocation routing policy to route requests to one of the six EC2 instance
- H. Create an Amazon CloudFront distributio
- I. Use the Route 53 record as the distribution's origin.
- J. Replace the two NLBs with two Application Load Balancers (ALBs). Create an Amazon Route 53 latency routing policy to route requests to one of the two ALB
- K. Create an Amazon CloudFront distributio
- L. Use the Route 53 record as the distribution's origin.

Answer: B

Explanation:

For standard accelerators, Global Accelerator uses the AWS global network to route traffic to the optimal regional endpoint based on health, client location, and policies that you configure, which increases the availability of your applications. Endpoints for standard accelerators can be Network Load Balancers, Application Load Balancers, Amazon EC2 instances, or Elastic IP addresses that are located in one AWS Region or multiple Regions.

<https://docs.aws.amazon.com/global-accelerator/latest/dg/what-is-global-accelerator.html>

NEW QUESTION 196

- (Topic 2)

A company is migrating an application from on-premises servers to Amazon EC2 instances. As part of the migration design requirements, a solutions architect must implement infrastructure metric alarms. The company does not need to take action if CPU utilization increases to more than 50% for a short burst of time. However, if the CPU

utilization increases to more than 50% and read IOPS on the disk are high at the same time, the company needs to act as soon as possible. The solutions architect also must reduce false alarms.

What should the solutions architect do to meet these requirements?

- A. Create Amazon CloudWatch composite alarms where possible.
- B. Create Amazon CloudWatch dashboards to visualize the metrics and react to issues quickly.
- C. Create Amazon CloudWatch Synthetics canaries to monitor the application and raise an alarm.
- D. Create single Amazon CloudWatch metric alarms with multiple metric thresholds where possible.

Answer: A

Explanation:

Composite alarms determine their states by monitoring the states of other alarms. You can **use composite alarms to reduce alarm noise**. For example, you can create a composite alarm where the underlying metric alarms go into ALARM when they meet specific conditions. You then can set up your composite alarm to go into ALARM and send you notifications when the underlying metric alarms go into ALARM by configuring the underlying metric alarms never to take actions.

Currently, composite alarms can take the following actions: https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/Create_Composite_Alarm.html

NEW QUESTION 201

- (Topic 3)

A company wants to migrate a Windows-based application from on premises to the AWS Cloud. The application has three tiers, a business tier, and a database tier with Microsoft SQL Server. The company wants to use specific features of SQL Server such as native backups and Data Quality Services. The company also needs to share files for process between the tiers.

How should a solution architect design the architecture to meet these requirements?

- A. Host all three on Amazon instance
- B. Use Mamazon FSx File Gateway for file sharing between tiers.
- C. Host all three on Amazon EC2 instance
- D. Use Amazon FSx for Windows file sharing between the tiers.
- E. Host the application tier and the business tier on Amazon EC2 instance
- F. Host the database tier on Amazon RD
- G. Use Amazon Elastic File system (Amazon EFS) for file sharing between the tiers.
- H. Host the application tier and the business tier on Amazon EC2 instance
- I. Host the database tier on Amazon RD
- J. Use a Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volume for file sharing between the tiers.

Answer: B

Explanation:

This solution will allow the company to host all three tiers on Amazon EC2 instances while using Amazon FSx for Windows File Server to provide Windows-based file sharing between the tiers. This will allow the company to use specific features of SQL Server, such as native backups and Data Quality Services, while sharing files for processing between the tiers.

NEW QUESTION 202

- (Topic 3)

A company has hundreds of Amazon EC2 Linux-based instances in the AWS Cloud. Systems administrators have used shared SSH keys to manage the instances. After a recent audit, the company's security team is mandating the removal of all shared keys. A solutions architect must design a solution that provides secure access to the EC2 instances.

Which solution will meet this requirement with the LEAST amount of administrative overhead?

- A. Use AWS Systems Manager Session Manager to connect to the EC2 instances.
- B. Use AWS Security Token Service (AWS STS) to generate one-time SSH keys on demand.
- C. Allow shared SSH access to a set of bastion instance
- D. Configure all other instances to allow only SSH access from the bastion instances
- E. Use an Amazon Cognito custom authorizer to authenticate user
- F. Invoke an AWS Lambda function to generate a temporary SSH key.

Answer: A

Explanation:

Session Manager is a fully managed AWS Systems Manager capability. With Session Manager, you can manage your Amazon Elastic Compute Cloud (Amazon EC2) instances, edge devices, on-premises servers, and virtual machines (VMs). You can use either an interactive one-click browser-based shell or the AWS Command Line Interface (AWS CLI). Session Manager provides secure and auditable node management without the need to open inbound ports, maintain bastion hosts, or manage SSH keys. Session Manager also allows you to comply with corporate policies that require controlled access to managed nodes, strict security practices, and fully auditable logs with node access details, while providing end users with simple one-click cross-platform access to your managed nodes.

<https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager.html>

NEW QUESTION 205

- (Topic 3)

A company is building a new web-based customer relationship management application. The application will use several Amazon EC2 instances that are backed by Amazon Elastic Block Store (Amazon EBS) volumes behind an Application Load Balancer (ALB). The application will also use an Amazon Aurora database. All data for the application must be encrypted at rest and in transit.

Which solution will meet these requirements?

- A. Use AWS Key Management Service (AWS KMS) certificates on the ALB to encrypt data in transit
- B. Use AWS Certificate Manager (ACM) to encrypt the EBS volumes and Aurora database storage at rest.
- C. Use the AWS root account to log in to the AWS Management Console
- D. Upload the company's encryption certificate
- E. While in the root account, select the option to turn on encryption for all data at rest and in transit for the account.
- F. Use a AWS Key Management Service (AWS KMS) to encrypt the EBS volumes and Aurora database storage at rest
- G. Attach an AWS Certificate Manager (ACM) certificate to the ALB to encrypt data in transit.
- H. Use BitLocker to encrypt all data at rest
- I. Import the company's TLS certificate keys to AWS Key Management Service (AWS KMS). Attach the KMS keys to the ALB to encrypt data in transit.

Answer: C

Explanation:

This option is the most efficient because it uses AWS Key Management Service (AWS KMS), which is a service that makes it easy for you to create and manage cryptographic keys and control their use across a wide range of AWS services and with your applications running on AWS. It also uses AWS KMS to encrypt the EBS volumes and Aurora database storage at rest, which provides data protection by encrypting your data with encryption keys that you manage. It also uses AWS Certificate Manager (ACM), which is a service that lets you easily provision, manage, and deploy public and private Secure Sockets Layer/Transport Layer Security (SSL/TLS) certificates for use with AWS services and your internal connected resources. It also attaches an ACM certificate to the ALB to encrypt data in transit, which provides data protection by enabling SSL/TLS encryption for connections between clients and the load balancer. This solution meets the requirement of encrypting all data for the application at rest and in transit. Option A is less efficient because it uses AWS KMS certificates on the ALB to encrypt data in transit, which is not possible as AWS KMS does not provide certificates but only keys. It also uses AWS Certificate Manager (ACM) to encrypt the EBS volumes and Aurora database storage at rest, which is not possible as ACM does not provide encryption but only certificates. Option B is less efficient because it uses the AWS root account to log in to the AWS Management Console, which is not recommended as it has unrestricted access to all resources in your account. It also uploads the company's encryption certificates, which is not necessary as ACM can provide certificates for free. It also selects the option to turn on encryption for all data at rest and in transit for the account, which is not possible as encryption settings are specific to each service and resource. Option D is less efficient because it uses BitLocker to encrypt all data at rest, which is a Windows feature that provides encryption for volumes on Windows servers. However, this does not provide encryption for Aurora database storage at rest, as Aurora runs on Linux servers. It also imports the company's TLS certificate keys to AWS KMS, which is not necessary as ACM can provide certificates for free. It also attaches the KMS keys to the ALB to encrypt data in transit, which is not possible as ALB requires certificates and not keys.

NEW QUESTION 206

- (Topic 3)

A company needs to migrate a legacy application from an on-premises data center to the AWS Cloud because of hardware capacity constraints. The application runs 24 hours a day, 7 days a week. The application database storage continues to grow over time.

What should a solution architect do to meet these requirements MOST cost-effectively?

- A. Migrate the application layer to Amazon EC2 Spot Instances Migrate the data storage layer to Amazon S3.
- B. Migrate the application layer to Amazon EC2 Reserved Instances Migrate the data storage layer to Amazon RDS On-Demand Instances.
- C. Migrate the application layer to Amazon EC2 Reserved instances Migrate the data storage layer to Amazon Aurora Reserved Instances.
- D. Migrate the application layer to Amazon EC2 On Demand Amazon Migrate the data storage layer to Amazon RDS Reserved instances.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.AuroraMySQL.html>

NEW QUESTION 209

- (Topic 3)

A company experienced a breach that affected several applications in its on-premises data center. The attacker took advantage of vulnerabilities in the custom

applications that were running on the servers. The company is now migrating its applications to run on Amazon EC2 instances. The company wants to implement a solution that actively scans for vulnerabilities on the EC2 instances and sends a report that details the findings. Which solution will meet these requirements?

- A. Deploy AWS Shield to scan the EC2 instances for vulnerabilities. Create an AWS Lambda function to log any findings to AWS CloudTrail.
- B. Deploy Amazon Macie and AWS Lambda functions to scan the EC2 instances for vulnerabilities. Log any findings to AWS CloudTrail.
- C. Turn on Amazon GuardDuty. Deploy the GuardDuty agents to the EC2 instances. Configure an AWS Lambda function to automate the generation and distribution of reports that detail the findings.
- D. Turn on Amazon Inspector. Deploy the Amazon Inspector agent to the EC2 instances. Configure an AWS Lambda function to automate the generation and distribution of reports that detail the findings.

Answer: D

Explanation:

Amazon Inspector:

- Performs active vulnerability scans of EC2 instances. It looks for software vulnerabilities, unintended network accessibility, and other security issues.
- Requires installing an agent on EC2 instances to perform scans. The agent must be deployed to each instance.
- Provides scheduled scan reports detailing any findings of security risks or vulnerabilities. These reports can be used to patch or remediate issues.
- Is best suited for proactively detecting security weaknesses and misconfigurations in your AWS environment.

NEW QUESTION 210

- (Topic 3)

A company has an application that runs on several Amazon EC2 instances. Each EC2 instance has multiple Amazon Elastic Block Store (Amazon EBS) data volumes attached to it. The application's EC2 instance configuration and data need to be backed up nightly. The application also needs to be recoverable in a different AWS Region.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. Write an AWS Lambda function that schedules nightly snapshots of the application's EBS volumes and copies the snapshots to a different Region.
- B. Create a backup plan by using AWS Backup to perform nightly backup.
- C. Copy the backups to another Region. Add the application's EC2 instances as resources.
- D. Create a backup plan by using AWS Backup to perform nightly backups. Copy the backups to another Region. Add the application's EBS volumes as resources.
- E. Write an AWS Lambda function that schedules nightly snapshots of the application's EBS volumes and copies the snapshots to a different Availability Zone.

Answer: B

Explanation:

The most operationally efficient solution to meet these requirements would be to create a backup plan by using AWS Backup to perform nightly backups and copying the backups to another Region. Adding the application's EBS volumes as resources will ensure that the application's EC2 instance configuration and data are backed up, and copying the backups to another Region will ensure that the application is recoverable in a different AWS Region.

NEW QUESTION 211

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

SAA-C03 Practice Exam Features:

- * SAA-C03 Questions and Answers Updated Frequently
- * SAA-C03 Practice Questions Verified by Expert Senior Certified Staff
- * SAA-C03 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * SAA-C03 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The SAA-C03 Practice Test Here](#)