Amazon-Web-Services

Exam Questions SAA-C03

AWS Certified Solutions Architect - Associate (SAA-C03)

NEW QUESTION 1
- (Topic 1)
A company observes an increase in Amazon EC2 costs in its most recent bill The billing team notices unwanted vertical scaling of instance types for a couple of EC2 instances A solutions architect needs to create a graph comparing the last 2 months of EC2 costs and perform an in-depth analysis to identify the root cause of the vertical scaling
How should the solutions architect generate the information with the LEAST operational overhead?

A. Use AWS Budgets to create a budget report and compare EC2 costs based on instance types
B. Use Cost Explorer's granular filtering feature to perform an in-depth analysis of EC2 costs based on instance types
C. Use graphs from the AWS Billing and Cost Management dashboard to compare EC2 costs based on instance types for the last 2 months
D. Use AWS Cost and Usage Reports to create a report and send it to an Amazon S3 bucket Use Amazon QuickSight with Amazon S3 as a source to generate an interactive graph based on instance types.

Answer: B

Explanation:
AWS Cost Explorer is a tool that enables you to view and analyze your costs and usage. You can explore your usage and costs using the main graph, the Cost Explorer cost and usage reports, or the Cost Explorer RI reports. You can view data for up to the last 12 months, forecast how much you're likely to spend for the next 12 months, and get recommendations for what Reserved Instances to purchase. You can use Cost Explorer to identify areas that need further inquiry and see trends that you can use to understand your
costs. https://docs.aws.amazon.com/cost-management/latest/userguide/ce-what-is.html

NEW QUESTION 2
- (Topic 1)
A company has an application that ingests incoming messages. These messages are then quickly consumed by dozens of other applications and microservices. The number of messages varies drastically and sometimes spikes as high as 100,000 each second. The company wants to decouple the solution and increase scalability.
Which solution meets these requirements?

A. Persist the messages to Amazon Kinesis Data Analytic
B. All the applications will read and process the messages.
C. Deploy the application on Amazon EC2 instances in an Auto Scaling group, which scales the number of EC2 instances based on CPU metrics.
D. Write the messages to Amazon Kinesis Data Streams with a single shar
E. All applications will read from the stream and process the messages.
F. Publish the messages to an Amazon Simple Notification Service (Amazon SNS) topic with one or more Amazon Simple Queue Service (Amazon SQS) subscription
G. All applications then process the messages from the queues.

Answer: D

Explanation:
https://aws.amazon.com/sqs/features/
By routing incoming requests to Amazon SQS, the company can decouple the job requests from the processing instances. This allows them to scale the number of instances based on the size of the queue, providing more resources when needed. Additionally, using an Auto Scaling group based on the queue size will automatically scale the number of instances up or down depending on the workload. Updating the software to read from the queue will allow it to process the job requests in a more efficient manner, improving the performance of the system.

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 poin
B. Use Amazon Inspector to scan me objects in the bucke
C. If objects contain Pl
D. trigger an S3 Lifecycle policy to remove the objects that contain Pll.
E. Use an Amazon S3 bucket as a secure transfer poin
F. Use Amazon Macie to scan the objects in the bucke
G. If objects contain Pl
H. Use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects mat contain Pll.
I. Implement custom scanning algorithms in an AWS Lambda functio
J. Trigger the function when objects are loaded into the bucke
K. It objects contain Rl
L. use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects that contain Pll.
M. Implement custom scanning algorithms in an AWS Lambda functio
N. Trigger the function when objects are loaded into the bucke
O. If objects contain Pl
P. use Amazon Simple Email Service (Amazon STS) to trigger a notification to the administrators and trigger on S3 Lifecycle policy to remove the objects mot 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 development team runs monthly resource-intensive tests on its general purpose Amazon RDS for MySQL DB instance with Performance Insights enabled. The testing lasts for 48 hours once a month and is the only process that uses the database. The team wants to reduce the cost of running the tests without reducing the compute and memory attributes of the DB instance.
Which solution meets these requirements MOST cost-effectively?

A. Stop the DB instance when tests are complete
B. Restart the DB instance when required.
C. Use an Auto Scaling policy with the DB instance to automatically scale when tests are completed.
D. Create a snapshot when tests are complete
E. Terminate the DB instance and restore the snapshot when required.
F. Modify the DB instance to a low-capacity instance when tests are complete
G. Modify the DB instance again when required.

Answer: A

Explanation:
To reduce the cost of running the tests without reducing the compute and memory attributes of the Amazon RDS for MySQL DB instance, the development team can stop the instance when tests are completed and restart it when required. Stopping the DB instance when not in use can help save costs because customers are only charged for storage while the DB instance is stopped. During this time, automated backups and automated DB instance maintenance are suspended. When the instance is restarted, it retains the same configurations, security groups, and DB parameter groups as when it was stopped.
Reference:
Amazon RDS Documentation: Stopping and Starting a DB instance (https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_StopInstance.html)

NEW QUESTION 5
- (Topic 1)
A company uses NFS to store large video files in on-premises network attached storage. Each video file ranges in size from 1MB to 500 GB. The total storage is 70 TB and is no longer growing. The company decides to migrate the video files to Amazon S3. The company must migrate the video files as soon as possible while using the least possible network bandwidth.
Which solution will meet these requirements?

A. Create an S3 bucket Create an IAM role that has permissions to write to the S3 bucke
B. Use the AWS CLI to copy all files locally to the S3 bucket.
C. Create an AWS Snowball Edge jo
D. Receive a Snowball Edge device on premise
E. Use the Snowball Edge client to transfer data to the devic
F. Return the device so that AWS can import the data into Amazon S3.
G. Deploy an S3 File Gateway on premise
H. Create a public service endpoint to connect to the S3 File Gateway Create an S3 bucket Create a new NFS file share on the S3 File Gateway Point the new file share to the S3 bucke
I. Transfer the data from the existing NFS file share to the S3 File Gateway.
J. Set up an AWS Direct Connect connection between the on-premises network and AW
K. Deploy an S3 File Gateway on premise
L. Create a public virtual interlace (VIF) to connect to the S3 File Gatewa
M. Create an S3 bucke
N. Create a new NFS file share on the S3 File Gatewa
O. Point the new file share to the S3 bucke
P. Transfer the data from the existing NFS file share to the S3 File Gateway.

Answer: B

Explanation:
The basic difference between Snowball and Snowball Edge is the capacity they provide. Snowball provides a total of 50 TB or 80 TB, out of which 42 TB or 72 TB is available, while Amazon Snowball Edge provides 100 TB, out of which 83 TB is available.

NEW QUESTION 6
- (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 7
- (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 8
- (Topic 1)
A company's containerized application runs on an Amazon EC2 instance. The application needs to download security certificates before it can communicate with other business applications. The company wants a highly secure solution to encrypt and decrypt the certificates in near real time. The solution also needs to store data in highly available storage after the data is encrypted.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create AWS Secrets Manager secrets for encrypted certificate
B. Manually update the certificates as neede
C. Control access to the data by using fine-grained IAM access.
D. Create an AWS Lambda function that uses the Python cryptography library to receive and perform encryption operation
E. Store the function in an Amazon S3 bucket.
F. Create an AWS Key Management Service (AWS KMS) customer managed ke
G. Allow the EC2 role to use the KMS key for encryption operation
H. Store the encrypted data on Amazon S3.
I. Create an AWS Key Management Service (AWS KMS) customer managed ke
J. Allow the EC2 role to use the KMS key for encryption operation
K. Store the encrypted data on Amazon Elastic Block Store (Amazon EBS) volumes.

Answer: D

NEW QUESTION 9
- (Topic 1)
A company has a website hosted on AWS. The website is behind an Application Load Balancer (ALB) that is configured to handle HTTP and HTTPS separately. The company wants to forward all requests to the website so that the requests will use HTTPS.
What should a solutions architect do to meet this requirement?

A. Update the ALB's network ACL to accept only HTTPS traffic
B. Create a rule that replaces the HTTP in the URL with HTTPS.
C. Create a listener rule on the ALB to redirect HTTP traffic to HTTPS.
D. Replace the ALB with a Network Load Balancer configured to use Server Name Indication (SNI).

Answer: C

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/elb-redirect-http-to-https- using-alb/
How can I redirect HTTP requests to HTTPS using an Application Load Balancer? Last updated: 2020-10-30 I want to redirect HTTP requests to HTTPS using Application Load Balancer listener rules. How can I do this? Resolution Reference: https://aws.amazon.com/premiumsupport/knowledge-center/elb-redirect-http-to-https- using-alb/

NEW QUESTION 10
- (Topic 1)
A company has created an image analysis application in which users can upload photos and add photo frames to their images. The users upload images and metadata to indicate which photo frames they want to add to their images. The application uses a single Amazon EC2 instance and Amazon DynamoDB to store

the metadata.
The application is becoming more popular, and the number of users is increasing. The company expects the number of concurrent users to vary significantly depending on the time of day and day of week. The company must ensure that the application can scale to meet the needs of the growing user base.
Which solution meats these requirements?

A. Use AWS Lambda to process the photo
B. Store the photos and metadata in DynamoDB.
C. Use Amazon Kinesis Data Firehose to process the photos and to store the photos and metadata.
D. Use AWS Lambda to process the photo
E. Store the photos in Amazon S3. Retain DynamoDB to store the metadata.
F. Increase the number of EC2 instances to thre
G. Use Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volumes to store the photos and metadata.

Answer: C

Explanation:
https://www.quora.com/How-can-I-use-DynamoDB-for-storing-metadata-for-Amazon-S3-objects
This solution meets the requirements of scalability, performance, and availability. AWS Lambda can process the photos in parallel and scale up or down automatically depending on the demand. Amazon S3 can store the photos and metadata reliably and durably, and provide high availability and low latency. DynamoDB can store the metadata efficiently and provide consistent performance. This solution also reduces the cost and complexity of managing EC2 instances and EBS volumes.
Option A is incorrect because storing the photos in DynamoDB is not a good practice, as it can increase the storage cost and limit the throughput. Option B is incorrect because Kinesis Data Firehose is not designed for processing photos, but for streaming data to destinations such as S3 or Redshift. Option D is incorrect because increasing the number of EC2 instances and using Provisioned IOPS SSD volumes does not guarantee scalability, as it depends on the load balancer and the application code. It also increases the cost and complexity of managing the infrastructure.
References:
? https://aws.amazon.com/certification/certified-solutions-architect-professional/
? https://www.examtopics.com/discussions/amazon/view/7193-exam-aws-certified-solutions-architect-professional-topic-1/
? https://aws.amazon.com/architecture/

NEW QUESTION 10
- (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 consol
D. Store the backup in an Amazon S3 bucke
E. Set an S3 Lifecycle configuration for the S3 bucket.
F. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda functio
G. Configure the Lambda function to back up the table and to store thebackup in an Amazon S3 bucke
H. Set an S3 Lifecycle configuration for the S3 bucket.

Answer: C

NEW QUESTION 13
- (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 distributio
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 16
- (Topic 1)
A company has applications that run on Amazon EC2 instances in a VPC. One of the applications needs to call the Amazon S3 API to store and read objects. According to the company's security regulations, no traffic from the applications is allowed to travel across the internet.
Which solution will meet these requirements?

A. Configure an S3 interface endpoint.
B. Configure an S3 gateway endpoint.
C. Create an S3 bucket in a private subnet.
D. Create an S3 bucket in the same Region as the EC2 instance.

Answer: B

Explanation:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/privatelink-interface-endpoints.html#types-of-vpc-endpoints-for-s3
https://docs.aws.amazon.com/vpc/latest/userguide/vpc-endpoints-s3.html

NEW QUESTION 19
- (Topic 1)
A company recently launched Linux-based application instances on Amazon EC2 in a private subnet and launched a Linux-based bastion host on an Amazon EC2 instance in a public subnet of a VPC A solutions architect needs to connect from the on-premises
network, through the company's internet connection to the bastion host and to the application servers The solutions architect must make sure that the security groups of all the EC2 instances will allow that access
Which combination of steps should the solutions architect take to meet these requirements? (Select TWO)

A. Replace the current security group of the bastion host with one that only allows inbound access from the application instances
B. Replace the current security group of the bastion host with one that only allows inbound access from the internal IP range for the company
C. Replace the current security group of the bastion host with one that only allows inbound access from the external IP range for the company
D. Replace the current security group of the application instances with one that allows inbound SSH access from only the private IP address of the bastion host
E. Replace the current security group of the application instances with one that allows inbound SSH access from only the public IP address of the bastion host

Answer: CD

Explanation:
https://digitalcloud.training/ssh-into-ec2-in-private-subnet/

NEW QUESTION 23
- (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 UR
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 nam
D. Point the alias record to the Regional API Gateway stage endpoin
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 endpoin
G. Associate the API Gateway endpoint with the company's domain nam
H. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM) in the same Regio
I. Attach the certificate to the API Gateway endpoin
J. Configure Route 53 to route traffic to the API Gateway endpoint.
K. Create a Regional API Gateway endpoin
L. Associate the API Gateway endpoint with the company's domain nam
M. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM) in the us-east-1 Regio
N. Attach the certificate to the API Gateway API
O. Create Route 53 DNS records with the company's domain nam
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 26
- (Topic 1)
A company is storing sensitive user information in an Amazon S3 bucket The company wants to provide secure access to this bucket from the application tier running on Ama2on EC2 instances inside a VPC.
Which combination of steps should a solutions architect take to accomplish this? (Select TWO.)

A. Configure a VPC gateway endpoint for Amazon S3 within the VPC
B. Create a bucket policy to make the objects to the S3 bucket public
C. Create a bucket policy that limits access to only the application tier running in the VPC
D. Create an IAM user with an S3 access policy and copy the IAM credentials to the EC2 instance
E. Create a NAT instance and have the EC2 instances use the NAT instance to access the S3 bucket

Answer: AC

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/s3-private-connection-no-authentication/

NEW QUESTION 28
- (Topic 1)
An Amazon EC2 administrator created the following policy associated with an IAM group containing several users

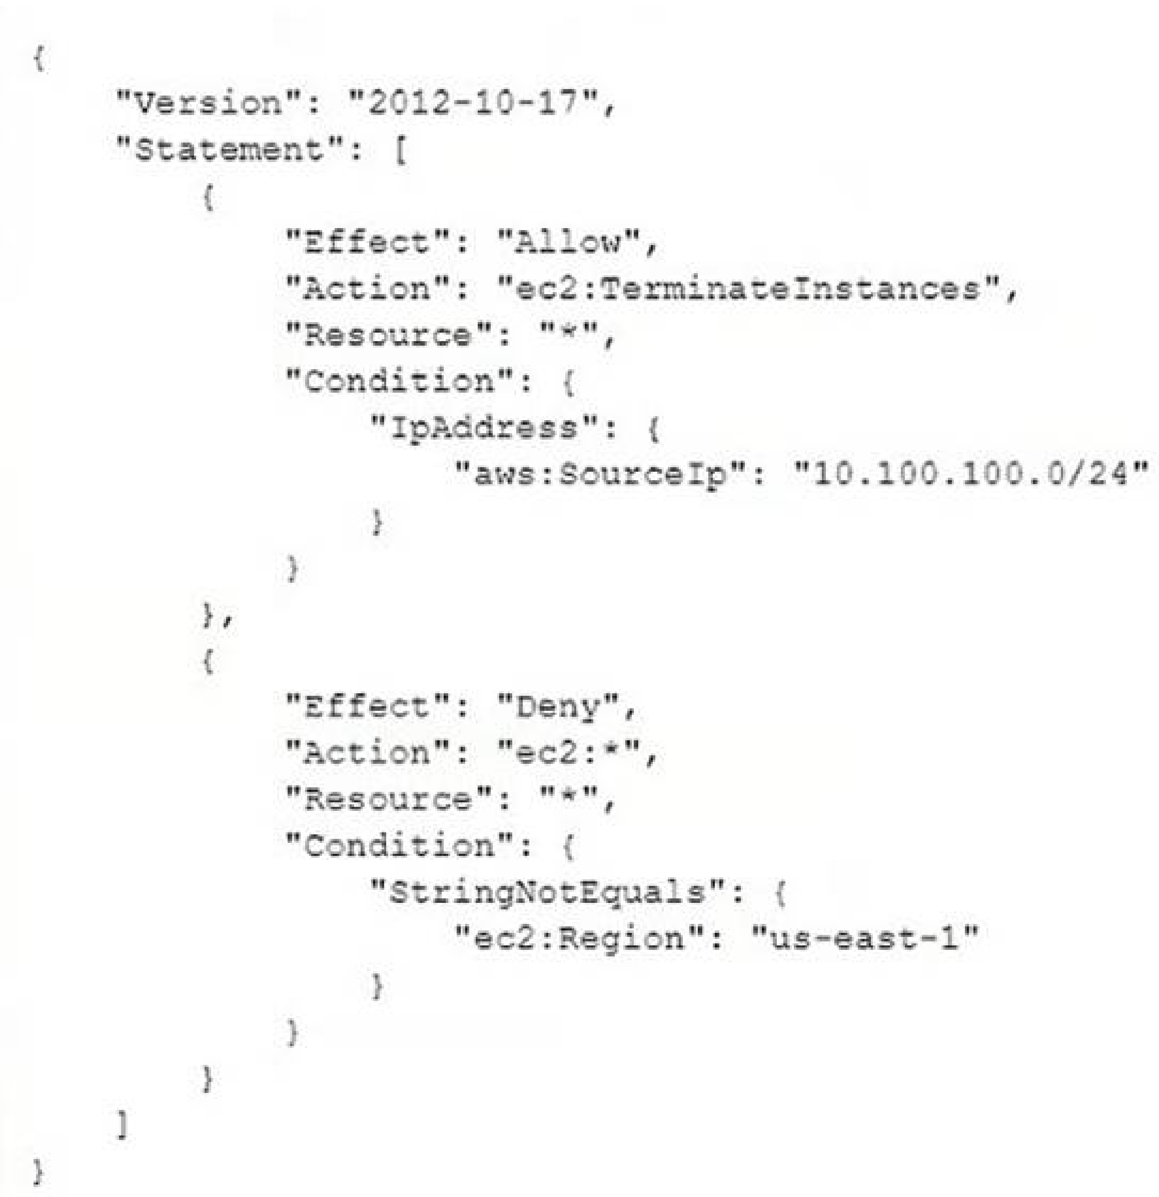

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "ec2:TerminateInstances",
            "Resource": "*",
            "Condition": {
                "IpAddress": {
                    "aws:SourceIp": "10.100.100.0/24"
                }
            }
        },
        {
            "Effect": "Deny",
            "Action": "ec2:*",
            "Resource": "*",
            "Condition": {
                "StringNotEquals": {
                    "ec2:Region": "us-east-1"
                }
            }
        }
    ]
}
```

What is the effect of this policy?

A. Users can terminate an EC2 instance in any AWS Region except us-east-1.
B. Users can terminate an EC2 instance with the IP address 10 100 100 1 in the us-east-1 Region
C. Users can terminate an EC2 instance in the us-east-1 Region when the user's source IP is 10.100.100.254.
D. Users cannot terminate an EC2 instance in the us-east-1 Region when the user's source IP is 10.100 100 254

Answer: C

Explanation:
as the policy prevents anyone from doing any EC2 action on any region except us-east-1 and allows only users with source ip 10.100.100.0/24 to terminate instances. So user with source ip 10.100.100.254 can terminate instances in us-east-1 region.

NEW QUESTION 29
- (Topic 1)
A social media company allows users to upload images to its website. The website runs on Amazon EC2 instances. During upload requests, the website resizes the images to a standard size and stores the resized images in Amazon S3. Users are experiencing slow upload requests to the website.
The company needs to reduce coupling within the application and improve website performance. A solutions architect must design the most operationally efficient process for image uploads.
Which combination of actions should the solutions architect take to meet these requirements? (Choose two.)

A. Configure the application to upload images to S3 Glacier.
B. Configure the web server to upload the original images to Amazon S3.
C. Configure the application to upload images directly from each user's browser to Amazon S3 through the use of a presigned URL.
D. Configure S3 Event Notifications to invoke an AWS Lambda function when an image is uploade
E. Use the function to resize the image
F. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that invokes an AWS Lambda function on a schedule to resize uploaded images.

Answer: CD

Explanation:
Amazon S3 is a highly scalable and durable object storage service that can store and retrieve any amount of data from anywhere on the web1. Users can configure the application to upload images directly from each user’s browser to Amazon S3 through the use of a presigned URL. A presigned URL is a URL that gives access to an object in an S3 bucket for a limited time and with a specific action, such as uploading an object2. Users can generate a presigned URL programmatically using the AWS SDKs or AWS CLI. By using a presigned URL, users can reduce coupling within the application and improve website performance, as they do not need to send the images to the web server first. AWS Lambda is a serverless compute service that runs code in response to events and automatically manages the underlying compute resources3. Users can configure S3 Event Notifications to invoke an AWS Lambda function when an image is uploaded. S3 Event Notifications is a feature that allows users to receive notifications when certain events happen in an S3 bucket, such as object creation or deletion. Users can configure S3 Event Notifications to invoke a Lambda function that resizes the image and stores it back in the same or a different S3 bucket. This way, users can offload the image resizing task from the web server to Lambda.

NEW QUESTION 30
- (Topic 1)
A company is deploying a new public web application to AWS. The application will run behind an Application Load Balancer (ALB). The application needs to be encrypted at the edge with an SSL/TLS certificate that is issued by an external certificate authority (CA). The certificate must be rotated each year before the certificate expires.

What should a solutions architect do to meet these requirements?

A. Use AWS Certificate Manager (ACM) to issue an SSL/TLS certificat
B. Apply the certificate to the AL
C. Use the managed renewal feature to automatically rotate the certificate.
D. Use AWS Certificate Manager (ACM) to issue an SSL/TLS certificat
E. Import the key material from the certificat
F. Apply the certificate to the AL
G. Use the managed renewal feature to automatically rotate the certificate.
H. Use AWS Certificate Manager (ACM) Private Certificate Authority to issue an SSL/TLS certificate from the root C
I. Apply the certificate to the AL
J. Use the managed renewal feature to automatically rotate the certificate.
K. Use AWS Certificate Manager (ACM) to import an SSL/TLS certificat
L. Apply the certificate to the AL
M. Use Amazon EventBridge (Amazon CloudWatch Events) to send a notification when the certificate is nearing expiratio
N. Rotate the certificate manually.

Answer: D

Explanation:
https://www.amazonaws.cn/en/certificate-manager/faqs/#Managed_renewal_and_deployment

NEW QUESTION 33
- (Topic 1)
A company has an AWS Glue extract. transform, and load (ETL) job that runs every day at the same time. The job processes XML data that is in an Amazon S3 bucket.
New data is added to the S3 bucket every day. A solutions architect notices that AWS Glue is processing all the data during each run.
What should the solutions architect do to prevent AWS Glue from reprocessing old data?

A. Edit the job to use job bookmarks.
B. Edit the job to delete data after the data is processed
C. Edit the job by setting the NumberOfWorkers field to 1.
D. Use a FindMatches machine learning (ML) transform.

Answer: C

Explanation:
This is the purpose of bookmarks: "AWS Glue tracks data that has already been processed during a previous run of an ETL job by persisting state information from the job run. This persisted state information is called a job bookmark. Job bookmarks help AWS Glue maintain state information and prevent the reprocessing of old data." https://docs.aws.amazon.com/glue/latest/dg/monitor-continuations.html

NEW QUESTION 37
- (Topic 1)
A company is preparing to store confidential data in Amazon S3 For compliance reasons the data must be encrypted at rest Encryption key usage must be logged tor auditing purposes. Keys must be rotated every year.
Which solution meets these requirements and «the MOST operationally efferent?

A. Server-side encryption with customer-provided keys (SSE-C)
B. Server-side encryption with Amazon S3 managed keys (SSE-S3)
C. Server-side encryption with AWS KMS (SSE-KMS) customer master keys (CMKs) with manual rotation
D. Server-side encryption with AWS KMS (SSE-KMS) customer master keys (CMKs) with automate rotation

Answer: D

Explanation:
https://docs.aws.amazon.com/kms/latest/developerguide/rotate-keys.html When you enable automatic key rotation for a customer managed key, AWS KMS generates new cryptographic material for the KMS key every year. AWS KMS also saves the KMS key's older cryptographic material in perpetuity so it can be used to decrypt data that the KMS key encrypted.
Key rotation in AWS KMS is a cryptographic best practice that is designed to be transparent and easy to use. AWS KMS supports optional automatic key rotation only for customer managed CMKs. Enable and disable key rotation. Automatic key rotation is disabled by default on customer managed CMKs. When you enable (or re-enable) key rotation, AWS KMS automatically rotates the CMK 365 days after the enable date and every 365 days thereafter.

NEW QUESTION 41
- (Topic 1)
A solutions architect is designing a two-tier web application The application consists of a public-facing web tier hosted on Amazon EC2 in public subnets The database tier consists of Microsoft SQL Server running on Amazon EC2 in a private subnet Security is a high priority for the company
How should security groups be configured in this situation? (Select TWO)

A. Configure the security group for the web tier to allow inbound traffic on port 443 from 0.0.0.0/0.
B. Configure the security group for the web tier to allow outbound traffic on port 443 from 0.0.0.0/0.
C. Configure the security group for the database tier to allow inbound traffic on port 1433 from the security group for the web tier.
D. Configure the security group for the database tier to allow outbound traffic on ports 443 and 1433 to the security group for the web tier.
E. Configure the security group for the database tier to allow inbound traffic on ports 443 and 1433 from the security group for the web tier.

Answer: AC

Explanation:
"Security groups create an outbound rule for every inbound rule." Not completely right. Statefull does NOT mean that if you create an inbound (or outbound) rule, it will create an outbound (or inbound) rule. What it does mean is: suppose you create an inbound rule on port 443 for the X ip. When a request enters on port 443 from X ip, it will allow traffic out for that request in the port 443. However, if you look at the outbound rules, there will not be any outbound rule on port 443 unless

explicitly create it. In ACLs, which are stateless, you would have to create an inbound rule to allow incoming requests and an outbound rule to allow your application responds to those incoming requests.
https://docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html#SecurityGro upRules

NEW QUESTION 45
- (Topic 1)
A company has an application that runs on Amazon EC2 instances and uses an Amazon
Aurora database. The EC2 instances connect to the database by using user names and passwords that are stored locally in a file. The company wants to minimize the operational overhead of credential management.
What should a solutions architect do to accomplish this goal?

A. Use AWS Secrets Manage
B. Turn on automatic rotation.
C. Use AWS Systems Manager Parameter Stor
D. Turn on automatic rotation.
E. Create an Amazon S3 bucket lo store objects that are encrypted with an AWS Key
F. Management Service (AWS KMS) encryption ke
G. Migrate the credential file to the S3 bucke
H. Point the application to the S3 bucket.
I. Create an encrypted Amazon Elastic Block Store (Amazon EBS) volume (or each EC2 instanc
J. Attach the new EBS volume to each EC2 instanc
K. Migrate the credential file to the new EBS volum
L. Point the application to the new EBS volume.

Answer: A

Explanation:
https://aws.amazon.com/cn/blogs/security/how-to-connect-to-aws-secrets-manager-service-within-a-virtual-private-cloud/
https://aws.amazon.com/blogs/security/rotate-amazon-rds-database-credentials-automatically-with-aws-secrets-manager/

NEW QUESTION 49
- (Topic 2)
A company has a Windows-based application that must be migrated to AWS. The application requires the use of a shared Windows file system attached to multiple Amazon EC2 Windows instances that are deployed across multiple Availability Zones.
What should a solutions architect do to meet this requirement?

A. Configure AWS Storage Gateway in volume gateway mod
B. Mount the volume to each Windows instance.
C. Configure Amazon FSx for Windows File Serve
D. Mount the Amazon FSx file system to each Windows instance.
E. Configure a file system by using Amazon Elastic File System (Amazon EFS). Mount the EFS file system to each Windows instance.
F. Configure an Amazon Elastic Block Store (Amazon EBS) volume with the required siz
G. Attach each EC2 instance to the volum
H. Mount the file system within the volume to each Windows instance.

Answer: B

Explanation:
This solution meets the requirement of migrating a Windows-based application that requires the use of a shared Windows file system attached to multiple Amazon EC2 Windows instances that are deployed across multiple Availability Zones. Amazon FSx for Windows File Server provides fully managed shared storage built on Windows Server, and delivers a wide range of data access, data management, and administrative capabilities. It supports the Server Message Block (SMB) protocol and can be mounted to EC2 Windows instances across multiple Availability Zones.
Option A is incorrect because AWS Storage Gateway in volume gateway mode provides cloud-backed storage volumes that can be mounted as iSCSI devices from on-premises application servers, but it does not support SMB protocol or EC2 Windows instances. Option C is incorrect because Amazon Elastic File System (Amazon EFS) provides a scalable and elastic NFS file system for Linux-based workloads, but it does not support SMB protocol or EC2 Windows instances. Option D is incorrect because Amazon Elastic Block Store (Amazon EBS) provides persistent block storage volumes for use with EC2 instances, but it does not support SMB protocol or attaching multiple instances to the same volume.
References:
? https://aws.amazon.com/fsx/windows/
? https://docs.aws.amazon.com/fsx/latest/WindowsGuide/using-file-shares.html

NEW QUESTION 50
- (Topic 2)
A company wants to direct its users to a backup static error page if the company's primary website is unavailable. The primary website's DNS records are hosted in Amazon Route 53. The domain is pointing to an Application Load Balancer (ALB). The company needs a solution that minimizes changes and infrastructure overhead.
Which solution will meet these requirements?

A. Update the Route 53 records to use a latency routing polic
B. Add a static error page that is hosted in an Amazon S3 bucket to the records so that the traffic is sent to the most responsive endpoints.
C. Set up a Route 53 active-passive failover configuratio
D. Direct traffic to a static error page that is hosted in an Amazon S3 bucket when Route 53 health checks determine that the ALB endpoint is unhealthy.
E. Set up a Route 53 active-active configuration with the ALB and an Amazon EC2 instance that hosts a static error page as endpoint
F. Configure Route 53 to send requests to the instance only if the health checks fail for the ALB.
G. Update the Route 53 records to use a multivalue answer routing polic
H. Create a health chec
I. Direct traffic to the website if the health check passe
J. Direct traffic to a static error page that is hosted in Amazon S3 if the health check does not pass.

Answer: B

Explanation:
This solution meets the requirements of directing users to a backup static error page if the primary website is unavailable, minimizing changes and infrastructure overhead. Route 53 active-passive failover configuration can route traffic to a primary resource when it is healthy or to a secondary resource when the primary resource is unhealthy. Route 53 health checks can monitor the health of the ALB endpoint and trigger the failover when needed. The static error page can be hosted in an S3 bucket that is configured as a website, which is a simple and cost-effective way to serve static content.
Option A is incorrect because using a latency routing policy can route traffic based on the lowest network latency for users, but it does not provide failover functionality. Option C is incorrect because using an active-active configuration with the ALB and an EC2 instance can increase the infrastructure overhead and complexity, and it does not guarantee that the EC2 instance will always be healthy. Option D is incorrect because using a multivalue answer routing policy can return multiple values for a query, but it does not provide failover functionality.
References:
? https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy- failover.html
? https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.html
? https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html

NEW QUESTION 54
- (Topic 2)
A company needs to save the results from a medical trial to an Amazon S3 repository. The repository must allow a few scientists to add new files and must restrict all other users to read-only access. No users can have the ability to modify or delete any files in the repository. The company must keep every file in the repository for a minimum of 1 year after its creation date.
Which solution will meet these requirements?

A. Use S3 Object Lock In governance mode with a legal hold of 1 year
B. Use S3 Object Lock in compliance mode with a retention period of 365 days.
C. Use an IAM role to restrict all users from deleting or changing objects in the S3 bucket Use an S3 bucket policy to only allow the IAM role
D. Configure the S3 bucket to invoke an AWS Lambda function every tune an object is added Configure the function to track the hash of the saved object to that modified objects can be marked accordingly

Answer: B

Explanation:
In compliance mode, a protected object version can't be overwritten or deleted by any user, including the root user in your AWS account. When an object is locked in compliance mode, its retention mode can't be changed, and its retention period can't be shortened. Compliance mode helps ensure that an object version can't be overwritten or deleted for the duration of the retention period. In governance mode, users can't overwrite or delete an object version or alter its lock settings unless they have special permissions. With governance mode, you protect objects against being deleted by most users, but you can still grant some users permission to alter the retention settings or delete the object if necessary. In Governance mode, Objects can be deleted by some users with special permissions, this is against the requirement.
Compliance:
- Object versions can't be overwritten or deleted by any user, including the root user
- Objects retention modes can't be changed, and retention periods can't be shortened
Governance:
- Most users can't overwrite or delete an object version or alter its lock settings
- Some users have special permissions to change the retention or delete the object

NEW QUESTION 59
- (Topic 2)
A company has a mulli-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 lo 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 tv/o 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 60
- (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 63
- (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 67
- (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 69
- (Topic 2)
A company is planning to move its data to an Amazon S3 bucket. The data must be encrypted when it is stored in the S3 bucket. Additionally, the encryption key must be automatically rotated every year.
Which solution will meet these requirements with the LEAST operational overhead?

A. Move the data to the S3 bucke
B. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Use the built-in key rotation behavior of SSE-S3 encryption keys.
C. Create an AWS Key Management Service {AWS KMS) customer managed ke
D. Enable automatic key rotatio
E. Set the S3 bucket's default encryption behavior to use the customer managed KMS ke
F. Move the data to the S3 bucket.
G. Create an AWS Key Management Service (AWS KMS) customer managed ke
H. Set the S3 bucket's default encryption behavior to use the customer managed KMS ke
I. Move the data to the S3 bucke
J. Manually rotate the KMS key every year.
K. Encrypt the data with customer key material before moving the data to the S3 bucke
L. Create an AWS Key Management Service (AWS KMS) key without key materia
M. Import the customer key material into the KMS ke
N. Enable automatic key rotation.

Answer: B

Explanation:
SSE-S3 - is free and uses AWS owned CMKs (CMK = Customer Master Key). The encryption key is owned and managed by AWS, and is shared among many accounts. Its rotation is automatic with time that varies as shown in the table here. The time is not explicitly defined.

SSE-KMS - has two flavors:
AWS managed CMK. This is free CMK generated only for your account. You can only view it policies and audit usage, but not manage it. Rotation is automatic - once per 1095 days (3 years),
Customer managed CMK. This uses your own key that you create and can manage. Rotation is not enabled by default. But if you enable it, it will be automatically rotated every 1 year. This variant can also use an imported key material by you. If you create such key with an imported material, there is no automated rotation. Only manual rotation.
SSE-C - customer provided key. The encryption key is fully managed by you outside of AWS. AWS will not rotate it.
This solution meets the requirements of moving data to an Amazon S3 bucket, encrypting the data when it is stored in the S3 bucket, and automatically rotating the encryption key every year with the least operational overhead. AWS Key Management Service (AWS KMS) is a service that enables you to create and manage encryption keys for your data. A customer managed key is a symmetric encryption key that you create and manage in AWS KMS. You can enable automatic key rotation for a customer managed key, which means that AWS KMS generates new cryptographic material for the key every year. You can set the S3 bucket's default encryption behavior to use the customer managed KMS key, which means that any object that is uploaded to the bucket without specifying an encryption method will be encrypted with that key.
Option A is incorrect because using server-side encryption with Amazon S3 managed encryption keys (SSE-S3) does not allow you to control or manage the encryption keys. SSE-S3 uses a unique key for each object, and encrypts that key with a master key that is regularly rotated by S3. However, you cannot enable or disable key rotation for SSE-S3 keys, or specify the rotation interval. Option C is incorrect because manually rotating the KMS key every year can increase the operational overhead and complexity, and it may not meet the requirement of rotating the key every year if you forget or delay the rotation
process. Option D is incorrect because encrypting the data with customer key material before moving the data to the S3 bucket can increase the operational overhead and complexity, and it may not provide consistent encryption for all objects in the bucket. Creating a KMS key without key material and importing the customer key material into the KMS key can enable you to use your own source of random bits to generate your KMS keys, but it does not support automatic key rotation.
References:
? https://docs.aws.amazon.com/kms/latest/developerguide/concepts.html
? https://docs.aws.amazon.com/kms/latest/developerguide/rotate-keys.html
? https://docs.aws.amazon.com/AmazonS3/latest/userguide/bucket-encryption.html

NEW QUESTION 70
- (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 AmazonEC2Ltnux 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 grou
B. Attach a single Provisioned IOPS SSD Amazon Elastic Block Store (Amazon E BS) volume to all the instances by using Amazon EBS Multi-Attach
C. Use a cluster placement grou
D. Create shared 'lie systems across the instances by using Amazon Elastic File System (Amazon EFS)
E. Use a partition placement grou
F. Create shared tile systems across the instances by using Amazon Elastic File System (Amazon EFS).
G. Use a spread placement grou
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-ebs-mult-attach

NEW QUESTION 74
- (Topic 2)
A company has an event-driven application that invokes AWS Lambda functions up to 800 times each minute with varying runtimes. The Lambda functions access data that is stored in an Amazon Aurora MySQL OB cluster. The company is noticing connection timeouts as user activity increases The database shows no signs of being overloaded. CPU. memory, and disk access metrics are all low.
Which solution will resolve this issue with the LEAST operational overhead?

A. Adjust the size of the Aurora MySQL nodes to handle more connection
B. Configure retry logic in the Lambda functions for attempts to connect to the database
C. Set up Amazon ElastiCache tor Redls to cache commonly read items from the databas
D. Configure the Lambda functions to connect to ElastiCache for reads.
E. Add an Aurora Replica as a reader nod
F. Configure the Lambda functions to connect to the reader endpoint of the OB cluster rather than lo the writer endpoint.
G. Use Amazon ROS Proxy to create a prox
H. Set the DB cluster as the target database Configure the Lambda functions lo connect to the proxy rather than to the DB cluster.

Answer: D

Explanation:
1. database shows no signs of being overloaded. CPU, memory, and disk access metrics are all low==>A and C out. We cannot only add nodes instance or add read replica, because database workload is totally fine, very low. 2. "least operational overhead"==>B out, because b need to configure lambda. 3. ROS proxy: Shares infrequently used connections; High availability with failover; Drives increased efficiency==>proxy can leverage failover to redirect traffic from timeout rds instance to
healthy rds instance. So D is right.

NEW QUESTION 78
- (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 80
- (Topic 2)
A company runs a web-based portal that provides users with global breaking news, local alerts, and weather updates. The portal delivers each user a personalized view by using mixture of static and dynamic content. Content is served over HTTPS through an API server running on an Amazon EC2 instance behind an Application Load Balancer (ALB). The company wants the portal to provide this content to its users across the world as quickly as possible.
How should a solutions architect design the application to ensure the LEAST amount of latency for all users?

A. Deploy the application stack in a single AWS Regio
B. Use Amazon CloudFront to serve all static and dynamic content by specifying the ALB as an origin.
C. Deploy the application stack in two AWS Region
D. Use an Amazon Route 53 latency routing policy to serve all content from the ALB in the closest Region.
E. Deploy the application stack in a single AWS Regio
F. Use Amazon CloudFront to serve the static conten
G. Serve the dynamic content directly from the ALB.
H. Deploy the application stack in two AWS Region
I. Use an Amazon Route 53 geolocation routing policy to serve all content from the ALB in the closest Region.

Answer: A

Explanation:
https://aws.amazon.com/blogs/networking-and-content-delivery/deliver-your-apps-dynamic-content-using-amazon-cloudfront-getting-started-template/

NEW QUESTION 84
- (Topic 2)
A company uses AWS Organizations to create dedicated AWS accounts for each business unit to manage each business unit's account independently upon request. The root email recipient missed a notification that was sent to the root user email address of one account. The company wants to ensure that all future notifications are not missed. Future notifications must be limited to account administrators.
Which solution will meet these requirements?

A. Configure the company's email server to forward notification email messages that are sent to the AWS account root user email address to all users in the organization.
B. Configure all AWS account root user email addresses as distribution lists that go to a few administrators who can respond to alert
C. Configure AWS account alternate contacts in the AWS Organizations console or programmatically.
D. Configure all AWS account root user email messages to be sent to one administrator who is responsible for monitoring alerts and forwarding those alerts to the appropriate groups.
E. Configure all existing AWS accounts and all newly created accounts to use the same root user email addres
F. Configure AWS account alternate contacts in the AWS Organizations console or programmatically.

Answer: B

Explanation:
Use a group email address for the management account's root user https://docs.aws.amazon.com/organizations/latest/userguide/orgs_best-practices_mgmt-acct.html#best-practices_mgmt-acct_email-address

NEW QUESTION 86
- (Topic 2)
A company runs its ecommerce application on AWS. Every new order is published as a message in a RabbitMQ queue that runs on an Amazon EC2 instance in a single Availability Zone. These messages are processed by a different application that runs on a separate EC2 instance. This application stores the details in a PostgreSQL database on another EC2 instance. All the EC2 instances are in the same Availability Zone.
The company needs to redesign its architecture to provide the highest availability with the least operational overhead.
What should a solutions architect do to meet these requirements?

A. Migrate the queue to a redundant pair (active/standby) of RabbitMQ instances on Amazon M
B. Create a Multi-AZ Auto Scaling group (or EC2 instances that host the applicatio
C. Create another Multi-AZAuto Scaling group for EC2 instances that host the PostgreSQL database.
D. Migrate the queue to a redundant pair (active/standby) of RabbitMQ instances on Amazon M
E. Create a Multi-AZ Auto Scaling group for EC2 instances that host the applicatio
F. Migrate the database to run on a Multi-AZ deployment of Amazon RDS for PostgreSQL.
G. Create a Multi-AZ Auto Scaling group for EC2 instances that host the RabbitMQ queu
H. Create another Multi-AZ Auto Scaling group for EC2 instances that host the application.Migrate the database to run on a Multi-AZ deployment of Amazon RDS fqjPostgreSQL.
I. Create a Multi-AZ Auto Scaling group for EC2 instances that host the RabbitMQ queue.Create another Multi-AZ Auto Scaling group for EC2 instances that host the applicatio
J. Create a third Multi-AZ Auto Scaling group for EC2 instances that host the PostgreSQL database.

Answer: B

Explanation:
Migrating to Amazon MQ reduces the overhead on the queue management. C and D are dismissed. Deciding between A and B means deciding to go for an AutoScaling group for EC2 or an RDS for Postgress (both multi- AZ). The RDS option has less operational impact, as provide as a service the tools and software required. Consider for instance, the effort to add an additional node like a read replica, to the DB. https://docs.aws.amazon.com/amazon-mq/latest/developer-guide/active-standby-broker- deployment.html https://aws.amazon.com/rds/postgresql/

NEW QUESTION 91
- (Topic 2)
A company runs a global web application on Amazon EC2 instances behind an Application Load Balancer The application stores data in Amazon Aurora. The company needs to create a disaster recovery solution and can tolerate up to 30 minutes of downtime and potential data loss. The solution does not need to handle the load when the primary infrastructure is healthy
What should a solutions architect do to meet these requirements?

A. Deploy the application with the required infrastructure elements in place Use Amazon Route 53 to configure active-passive failover Create an Aurora Replica in a second AWS Region
B. Host a scaled-down deployment of the application in a second AWS Region Use Amazon Route 53 to configure active-active failover Create an Aurora Replica in the second Region
C. Replicate the primary infrastructure in a second AWS Region Use Amazon Route 53 to configure active-active failover Create an Aurora database that is restored from the latest snapshot
D. Back up data with AWS Backup Use the backup to create the required infrastructure in a second AWS Region Use Amazon Route 53 to configure active-passive failover Create an Aurora second primary instance in the second Region

Answer: A

Explanation:
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-types.html

NEW QUESTION 95
- (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 Manage
B. Create an IAM role for the polic
C. Update the trust relationship of the rol
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 stoppe
F. Invalidate the cache after the DB instance is started.
G. Launch an Amazon EC2 instanc
H. Create an IAM role that grants access to Amazon RD
I. Attach the role to the EC2 instanc
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 instanc
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 96
- (Topic 2)
A company needs to move data from an Amazon EC2 instance to an Amazon S3 bucket. The company must ensure that no API calls and no data are routed through public internet routes. Only the EC2 instance can have access to upload data to the S3 bucket.
Which solution will meet these requirements?

A. Create an interface VPC endpoint for Amazon S3 in the subnet where the EC2 instance is locate
B. Attach a resource policy to the S3 bucket to only allow the EC2 instance's IAM role for access.
C. Create a gateway VPC endpoint for Amazon S3 in the Availability Zone where the EC2 instance is locate
D. Attach appropriate security groups to the endpoin
E. Attach a resource policy lo the S3 bucket to only allow the EC2 instance's IAM role for access.
F. Run the nslookup tool from inside the EC2 instance to obtain the private IP address of the S3 bucket's service API endpoin
G. Create a route in the VPC route table to provide theEC2 instance with access to the S3 bucke
H. Attach a resource policy to the S3 bucket to only allow the EC2 instance's IAM role for access.
I. Use the AWS provided, publicly available ip-ranges.json tile to obtain the private IP address of the S3 bucket's service API endpoin
J. Create a route in the VPC route table to provide the EC2 instance with access to the S3 bucke
K. Attach a resource policy to the S3 bucket to only allow the EC2 instance's IAM role for access.

Answer: A

Explanation:

(https://aws.amazon.com/blogs/security/how-to-restrict-amazon-s3-bucket-access-to-a-specific-iam-role/)

NEW QUESTION 100
- (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 105
- (Topic 2)
A company is running a multi-tier web application on premises. The web application is containerized and runs on a number of Linux hosts connected to a PostgreSQL database that contains user records. The operational overhead of maintaining the infrastructure and capacity planning is limiting the company's growth. A solutions architect must improve the application's infrastructure.
Which combination of actions should the solutions architect take to accomplish this? (Choose two.)

A. Migrate the PostgreSQL database to Amazon Aurora
B. Migrate the web application to be hosted on Amazon EC2 instances.
C. Set up an Amazon CloudFront distribution for the web application content.
D. Set up Amazon ElastiCache between the web application and the PostgreSQL database.
E. Migrate the web application to be hosted on AWS Fargate with Amazon Elastic Container Service (Amazon ECS).

Answer: AE

Explanation:
Amazon Aurora is a fully managed, scalable, and highly available relational database service that is compatible with PostgreSQL. Migrating the database to Amazon Aurora would reduce the operational overhead of maintaining the database infrastructure and allow the company to focus on building and scaling the application. AWS Fargate is a fully managed container orchestration service that enables users to run containers without the need to manage the underlying EC2 instances. By using AWS Fargate with Amazon Elastic Container Service (Amazon ECS), the solutions architect can improve the scalability and efficiency of the web application and reduce the operational overhead of maintaining the underlying infrastructure.

NEW QUESTION 106
- (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_Al arm.html

NEW QUESTION 108
- (Topic 3)
A company is migrating an old application to AWS The application runs a batch job every hour and is CPU intensive The batch job takes 15 minutes on average with an on-premises server The server has 64 virtual CPU (vCPU) and 512 GiB of memory
Which solution will run the batch job within 15 minutes with the LEAST operational overhead?

A. Use AWS Lambda with functional scaling
B. Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate
C. Use Amazon Lightsail with AWS Auto Scaling
D. Use AWS Batch on Amazon EC2

Answer: D

Explanation:
Use AWS Batch on Amazon EC2. AWS Batch is a fully managed batch processing service that can be used to easily run batch jobs on Amazon EC2 instances. It can scale the number of instances to match the workload, allowing the batch job to be completed in the desired time frame with minimal operational overhead.
Using AWS Lambda with Amazon API Gateway - AWS Lambda https://docs.aws.amazon.com/lambda/latest/dg/services-apigateway.html

AWS Lambda FAQs https://aws.amazon.com/lambda/faqs/

NEW QUESTION 112
- (Topic 3)
A media company hosts its website on AWS. The website application's architecture includes a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB) and a database that is hosted on Amazon Aurora The company's cyber security teem reports that the application is vulnerable to SOL injection.
How should the company resolve this issue?

A. Use AWS WAF in front of the ALB Associate the appropriate web ACLs with AWS WAF.
B. Create an ALB listener rule to reply to SQL injection with a fixed response
C. Subscribe to AWS Shield Advanced to block all SQL injection attempts automatically.
D. Set up Amazon Inspector to block all SOL injection attempts automatically

Answer: A

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/waf-block-common-attacks/#:~:text=To%20protect%20your%20applications%20against,%2C%20query%20string%2C%20or%20URI. --
----------------------------- Protect against SQL injection and cross-site scripting To protect your applications against SQL injection and cross-site scripting (XSS) attacks, use the built- in SQL injection and cross-site scripting engines. Remember that attacks can be performed on different parts of the HTTP request, such as the HTTP header, query string, or URI. Configure the AWS WAF rules to inspect different parts of the HTTP request against the built-in mitigation engines.

NEW QUESTION 115
- (Topic 3)
A company hosts a marketing website in an on-premises data center. The website consists of static documents and runs on a single server. An administrator updates the website content infrequently and uses an SFTP client to upload new documents.
The company decides to host its website on AWS and to use Amazon CloudFront. The company's solutions architect creates a CloudFront distribution. The solutions architect
must design the most cost-effective and resilient architecture for website hosting to serve as the CloudFront origin.
Which solution will meet these requirements?

A. Create a virtual server by using Amazon Lightsai
B. Configure the web server in the Lightsail instanc
C. Upload website content by using an SFTP client.
D. Create an AWS Auto Scaling group for Amazon EC2 instance
E. Use an Application Load Balance
F. Upload website content by using an SFTP client.
G. Create a private Amazon S3 bucke
H. Use an S3 bucket policy to allow access from a CloudFront origin access identity (OAI). Upload website content by using theAWSCLI.
I. Create a public Amazon S3 bucke
J. Configure AWS Transfer for SFT
K. Configure the S3 bucket for website hostin
L. Upload website content by using the SFTP client.

Answer: C

Explanation:
https://docs.aws.amazon.com/cli/latest/reference/transfer/describe- server.html

NEW QUESTION 117
- (Topic 3)
A solutions architect is designing the architecture for a software demonstration environment The environment will run on Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer (ALB) The system will experience significant increases in traffic during working hours but Is not required to operate on weekends.
Which combination of actions should the solutions architect take to ensure that the system can scale to meet demand? (Select TWO)

A. Use AWS Auto Scaling to adjust the ALB capacity based on request rate
B. Use AWS Auto Scaling to scale the capacity of the VPC internet gateway
C. Launch the EC2 instances in multiple AWS Regions to distribute the load across Regions
D. Use a target tracking scaling policy to scale the Auto Scaling group based on instance CPU utilization
E. Use scheduled scaling to change the Auto Scaling group minimum, maximum, and desired capacity to zero for weekends Revert to the default values at the start of the week

Answer: DE

Explanation:
https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-scaling-target- tracking.html#target-tracking-choose-metrics
A target tracking scaling policy is a type of dynamic scaling policy that adjusts the capacity of an Auto Scaling group based on a specified metric and a target value1. A target tracking scaling policy can automatically scale out or scale in the Auto Scaling group to keep the actual metric value at or near the target value1. A target tracking scaling policy is suitable for scenarios where the load on the application changes frequently and unpredictably, such as during working hours2.
To meet the requirements of the scenario, the solutions architect should use a target tracking scaling policy to scale the Auto Scaling group based on instance CPU utilization. Instance CPU utilization is a common metric that reflects the demand on the
application1. The solutions architect should specify a target value that represents the ideal average CPU utilization level for the application, such as 50 percent1. Then, the Auto Scaling group will scale out or scale in to maintain that level of CPU utilization.
Scheduled scaling is a type of scaling policy that performs scaling actions based on a date and time3. Scheduled scaling is suitable for scenarios where the load on the application changes periodically and predictably, such as on weekends2.
To meet the requirements of the scenario, the solutions architect should also use scheduled scaling to change the Auto Scaling group minimum, maximum, and desired capacity to zero for weekends. This way, the Auto Scaling group will terminate all instances on weekends when they are not required to operate. The solutions architect should also revert to the default values at the start of the week, so that the Auto Scaling group can resume normal operation.

NEW QUESTION 121
- (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. & days a week,. The application database storage continues to grow over time.
What should a solution architect do to meet these requirements MOST cost-affectivity?

A. Migrate the application layer to Amazon FC2 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 126
- (Topic 3)
A company wants to run an in-memory database for a latency-sensitive application that runs on Amazon EC2 instances. The application processes more than 100,000 transactions each minute and requires high network throughput. A solutions architect needs to provide a cost-effective network design that minimizes data transfer charges.
Which solution meets these requirements?

A. Launch all EC2 instances in the same Availability Zone within the same AWS Regio
B. Specify a placement group with cluster strategy when launching EC2 instances.
C. Launch all EC2 instances in different Availability Zones within the same AWS Regio
D. Specify a placement group with partition strategy when launching EC2 instances.
E. Deploy an Auto Scaling group to launch EC2 instances in different Availability Zones based on a network utilization target.
F. Deploy an Auto Scaling group with a step scaling policy to launch EC2 instances in different Availability Zones.

Answer: A

Explanation:
• Launching instances within a single AZ and using a cluster placement group provides the lowest network latency and highest bandwidth between instances. This maximizes performance for an in-memory database and high-throughput application.
• Communications between instances in the same AZ and placement group are free, minimizing data transfer charges. Inter-AZ and public IP traffic can incur charges.
• A cluster placement group enables the instances to be placed close together within the
AZ, allowing the high network throughput required. Partition groups span AZs, reducing bandwidth.
• Auto Scaling across zones could launch instances in AZs that increase data transfer charges. It may reduce network throughput, impacting performance.

NEW QUESTION 129
- (Topic 3)
A company is using Amazon Route 53 latency-based routing to route requests to its UDP- based application for users around the world. The application is hosted on redundant servers in the company's on-premises data centers in the United States. Asia, and Europe. The company's compliance requirements state that the application must be hosted on premises The company wants to improve the performance and availability of the application
What should a solutions architect do to meet these requirements?

A. A Configure three Network Load Balancers (NLBs) in the three AWS Regions to address the on-premises endpoints Create an accelerator by using AWS Global Accelerator, and register the NLBs as its endpoint
B. Provide access to the application by using a CNAME that points to the accelerator DNS
C. Configure three Application Load Balancers (ALBs) in the three AWS Regions to address the on-premises endpoint
D. Create an accelerator by using AWS Global Accelerator and register the ALBs as its endpoints Provide access to the application by using a CNAME that points to the accelerator DNS
E. Configure three Network Load Balancers (NLBs) in the three AWS Regions to address the on-premises endpoints In Route 53. create a latency-based record that points to the three NLB
F. and use it as an origin for an Amazon CloudFront distribution Provide access to the application by using a CNAME that points to the CloudFront DNS
G. Configure three Application Load Balancers (ALBs) in the three AWS Regions to address the on-premises endpoints In Route 53 create a latency-based record that points to the three ALBs and use it as an origin for an Amazon CloudFront distribution- Provide access to the application by using a CNAME that points to the CloudFront DNS

Answer: A

Explanation:
https://aws.amazon.com/step-functions/#:~:text=AWS%20Step%20Functions%20is%20a,machine%20learning%20(ML)%20pipelines.
"A common use case for AWS Step Functions is a task that requires human intervention (for example, an approval process). Step Functions makes it easy to coordinate the components of distributed applications as a series of steps in a visual workflow called a state machine. You can quickly build and run state machines to execute the steps of your application in a reliable and scalable fashion. (https://aws.amazon.com/pt/blogs/compute/implementing-serverless-manual-approval- steps-in-aws-step-functions-and-amazon-api-gateway/)"

NEW QUESTION 130
- (Topic 3)
A developer has an application that uses an AWS Lambda function to upload files to Amazon S3 and needs the required permissions to perform the task The developer already has an IAM user with valid IAM credentials required for Amazon S3
What should a solutions architect do to grant the permissions?

A. Add required IAM permissions in the resource policy of the Lambda function
B. Create a signed request using the existing IAM credentials n the Lambda function
C. Create a new IAM user and use the existing IAM credentials in the Lambda function.

D. Create an IAM execution role with the required permissions and attach the IAM rote to the Lambda function

Answer: D

Explanation:
To grant the necessary permissions to an AWS Lambda function to upload files to Amazon S3, a solutions architect should create an IAM execution role with the required permissions and attach the IAM role to the Lambda function. This approach follows the principle of least privilege and ensures that the Lambda function can only access the resources it needs to perform its specific task.

NEW QUESTION 133
- (Topic 3)
A company has an application thai 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 137
- (Topic 3)
A company runs a public three-Tier web application in a VPC The application runs on Amazon EC2 instances across multiple Availability Zones. The EC2 instances that run in private subnets need to communicate with a license server over the internet The company needs a managed solution that minimizes operational maintenance
Which solution meets these requirements"

A. Provision a NAT instance in a public subnet Modify each private subnets route table with a default route that points to the NAT instance
B. Provision a NAT instance in a private subnet Modify each private subnet's route table with a default route that points to the NAT instance
C. Provision a NAT gateway in a public subnet Modify each private subnet's route table with a default route that points to the NAT gateway
D. Provision a NAT gateway in a private subnet Modify each private subnet's route table with a default route that points to the NAT gateway

Answer: C

Explanation:
A NAT gateway is a type of network address translation (NAT) device that enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating connections with those instances1. A NAT gateway is a managed service that requires minimal operational maintenance and can handle up to 45 Gbps of bursty traffic1. A NAT gateway is suitable for scenarios where EC2 instances in private subnets need to communicate with a license server over the internet, such as the three-tier web application in the scenario1.
To meet the requirements of the scenario, the solutions architect should provision a NAT gateway in a public subnet. The solutions architect should also modify each private subnet’s route table with a default route that points to the NAT gateway1. This way, the EC2 instances that run in private subnets can access the license server over the internet through the NAT gateway.

NEW QUESTION 141
- (Topic 3)
A company is hosting a three-tier ecommerce application in the AWS Cloud. The company hosts the website on Amazon S3 and integrates the website with an API that handles sales requests. The company hosts the API on three Amazon EC2 instances behind an Application Load Balancer (ALB). The API consists of static and dynamic front-end content along with backend workers that process sales requests asynchronously.
The company is expecting a significant and sudden increase in the number of sales requests during events for the launch of new products
What should a solutions architect recommend to ensure that all the requests are processed successfully?

A. Add an Amazon CloudFront distribution for the dynamic conten
B. Increase the number of EC2 instances to handle the increase in traffic.
C. Add an Amazon CloudFront distribution for the static conten
D. Place the EC2 instances in an Auto Scaling group to launch new instances based on network traffic.
E. Add an Amazon CloudFront distribution for the dynamic conten
F. Add an Amazon ElastiCache instance in front of the ALB to reduce traffic for the API to handle.
G. Add an Amazon CloudFront distribution for the static conten
H. Add an Amazon Simple Queue Service (Amazon SOS) queue to receive requests from the website for later processing by the EC2 instances.

Answer: B

Explanation:
This option is the most efficient because it uses Amazon CloudFront, which is a web service that speeds up distribution of your static and dynamic web content, such as .html,
.css, .js, and image files, to your users1. It also uses a CloudFront distribution for the static content, which reduces the load on the EC2 instances and improves the performance and availability of the website. It also uses an Auto Scaling group to launch new instances based on network traffic, which automatically adjusts the compute capacity of your EC2 instances based on load or a schedule2. This solution meets the requirement of ensuring that all the requests are processed successfully during events for the launch of new products. Option A is less efficient because it uses a CloudFront distribution for the dynamic content, which is not necessary as the dynamic content is already handled by the API on the EC2 instances. It also increases the number of EC2 instances to handle the increase in traffic, which could incur higher costs and complexity than using an Auto Scaling group. Option C is less efficient because it uses an Amazon ElastiCache instance

in front of the ALB to reduce traffic for the API to handle, which is a way to provide a fully managed in-memory data store service that provides sub-millisecond latency for caching and data processing3. However, this could introduce additional complexity and latency, and does not scale automatically based on network traffic. Option D is less efficient because it uses an Amazon Simple Queue Service (Amazon SQS) queue to receive requests from the website for later processing by the EC2 instances, which is a way to send, store, and receive messages between software components at any volume. However, this does not provide a faster response time to the users as they have to wait for their requests to be processed by the EC2 instances.

NEW QUESTION 145

- (Topic 3)

A solutions architect observes that a nightly batch processing job is automatically scaled up for 1 hour before the desired Amazon EC2 capacity is reached. The peak capacity is the ‘same every night and the batch jobs always start at 1 AM. The solutions architect needs to find a cost-effective solution that will allow for the desired EC2 capacity to be reached quickly and allow the Auto Scaling group to scale down after the batch jobs are complete.
What should the solutions architect do to meet these requirements?

A. Increase the minimum capacity for the Auto Scaling group.
B. Increase the maximum capacity for the Auto Scaling group.
C. Configure scheduled scaling to scale up to the desired compute level.
D. Change the scaling policy to add more EC2 instances during each scaling operation.

Answer: C

Explanation:
By configuring scheduled scaling, the solutions architect can set the Auto Scaling group to automatically scale up to the desired compute level at a specific time (IAM) when the batch job starts and then automatically scale down after the job is complete. This will allow the desired EC2 capacity to be reached quickly and also help in reducing the cost.

NEW QUESTION 147

- (Topic 3)

A company wants to restrict access to the content of one of its man web applications and to protect the content by using authorization techniques available on AWS. The company wants to implement a serverless architecture end an authentication solution for fewer tian 100 users. The solution needs to integrate with the main web application and serve web content globally. The solution must also scale as to company's user base grows while providing lowest login latency possible.

A. Use Amazon Cognito tor authenticatio
B. Use Lambda#Edge tor authorization Use Amazon CloudFront 10 serve the web application globally
C. Use AWS Directory Service for Microsoft Active Directory tor authentication Use AWS Lambda for authorization Use an Application Load Balancer to serve the web application globally
D. Usa Amazon Cognito for authentication Use AWS Lambda tor authorization Use Amazon S3 Transfer Acceleration 10 serve the web application globally.
E. Use AWS Directory Service for Microsoft Active Directory for authentication Use Lambda@Edge for authorization Use AWS Elastic Beanstalk to serve the web application.

Answer: A

Explanation:
https://aws.amazon.com/blogs/networking-and-content-delivery/adding-http- security-headers-using-lambdaedge-and-amazon-cloudfront/
Amazon CloudFront is a global content delivery network (CDN) service that can securely deliver web content, videos, and APIs at scale. It integrates with Cognito for authentication and with Lambda@Edge for authorization, making it an ideal choice for serving web content globally. Lambda@Edge is a service that lets you run AWS Lambda functions globally closer to users, providing lower latency and faster response times. It can also handle authorization logic at the edge to secure content in CloudFront. For this scenario, Lambda@Edge can provide authorization for the web application while leveraging the low- latency benefit of running at the edge.

NEW QUESTION 148

- (Topic 3)

A company wants to create an application to store employee data in a hierarchical structured relationship. The company needs a minimum-latency response to high-traffic queries for the employee data and must protect any sensitive data. The company also need to receive monthly email messages if any financial information is present in the employee data.
Which combination of steps should a solutin architect take to meet these requirement? (
Select TWO.)

A. Use Amazon Redshift to store the employee data in hierarchie
B. Unload the data to Amazon S3 every month.
C. Use Amazon DynamoDB to store the employee data in hierarchies Export the data to Amazon S3 every month.
D. Configure Amazon Macie for the AWS account Integrate Macie with Amazon EventBridge to send monthly events to AWS Lambda.
E. Use Amazon Athena to analyze the employee data in Amazon S3 integrate Athena with Amazon QuickSight to publish analysis dashboards and share the dashboards with users.
F. Configure Amazon Macie for the AWS accoun
G. integrate Macie with Amazon EventBridge to send monthly notifications through an Amazon Simple Notification Service (Amazon SNS) subscription.

Answer: BE

Explanation:
https://docs.aws.amazon.com/prescriptive-guidance/latest/dynamodb-hierarchical-data-model/introduction.html

NEW QUESTION 150

- (Topic 3)

A rapidly growing ecommerce company is running its workloads in a single AWS Region. A solutions architect must create a disaster recovery (DR) strategy that includes a different AWS Region The company wants its database to be up to date in the DR Region with the least possible latency The remaining infrastructure in the DR Region needs to run at reduced capacity and must be able to scale up it necessary
Which solution will meet these requirements with the LOWEST recovery time objective (RTO)?

A. Use an Amazon Aurora global database with a pilot light deployment

B. Use an Amazon Aurora global database with a warm standby deployment
C. Use an Amazon RDS Multi-AZ DB instance with a pilot light deployment
D. Use an Amazon RDS Multi-AZ DB instance with a warm standby deployment

Answer: B

Explanation:
https://docs.aws.amazon.com/whitepapers/latest/disaster-recovery-workloads-on-aws/disaster-recovery-options-in-the-cloud.html

NEW QUESTION 153
- (Topic 3)
A company uses a payment processing system that requires messages for a particular payment ID to be received in the same order that they were sent Otherwise, the payments might be processed incorrectly.
Which actions should a solutions architect take to meet this requirement? (Select TWO.)

A. Write the messages to an Amazon DynamoDB table with the payment ID as the partition key
B. Write the messages to an Amazon Kinesis data stream with the payment ID as the partition key.
C. Write the messages to an Amazon ElastiCache for Memcached cluster with the payment ID as the key
D. Write the messages to an Amazon Simple Queue Service (Amazon SQS) queue Set the message attribute to use the payment ID
E. Write the messages to an Amazon Simple Queue Service (Amazon SQS) FIFO queu
F. Set the message group to use the payment ID.

Answer: BE

Explanation:
1) SQS FIFO queues guarantee that messages are received in the exact order they are sent. Using the payment ID as the message group ensures all messages for a payment ID are received sequentially. 2) Kinesis data streams can also enforce ordering on a per partition key basis. Using the payment ID as the partition key will ensure strict ordering of messages for each payment ID.

NEW QUESTION 158
- (Topic 3)
A meteorological startup company has a custom web application to sell weather data to its users online. The company uses Amazon DynamoDB to store is data and wants to bu4d a new service that sends an alert to the managers of four Internal teams every time a new weather event is recorded. The company does not want true new service to affect the performance of the current application
What should a solutions architect do to meet these requirement with the LEAST amount of operational overhead?

A. Use DynamoDB transactions to write new event data to the table Configure the transactions to notify internal teams.
B. Have the current application publish a message to four Amazon Simple Notification Service (Amazon SNS) topic
C. Have each team subscribe to one topic.
D. Enable Amazon DynamoDB Streams on the tabl
E. Use triggers to write to a mingle Amazon Simple Notification Service (Amazon SNS) topic to which the teams can subscribe.
F. Add a custom attribute to each record to flag new item
G. Write a cron job that scans the table every minute for items that are new and notifies an Amazon Simple Queue Service (Amazon SOS) queue to which the teams can subscribe.

Answer: C

Explanation:
The best solution to meet these requirements with the least amount of operational overhead is to enable Amazon DynamoDB Streams on the table and use triggers to write to a single Amazon Simple Notification Service (Amazon SNS) topic to which the teams can subscribe. This solution requires minimal configuration and infrastructure setup, and Amazon DynamoDB Streams provide a low-latency way to capture changes to the DynamoDB table. The triggers automatically capture the changes and publish them to the SNS topic, which notifies the internal teams.

NEW QUESTION 160
- (Topic 3)
A company has an on-premises MySQL database used by the global tales team with infrequent access patterns. The sales team requires the database to have minimal downtime. A database administrate wants to migrate this database to AWS without selecting a particular instance type in anticipation of more users In the future.
Which service should a solutions architect recommend?

A. Amazon Aurora MySQL
B. Amazon Aurora Serverless tor MySQL
C. Amazon Redshift Spectrum
D. Amazon RDS for MySQL

Answer: B

Explanation:
Amazon Aurora Serverless for MySQL is a fully managed, auto-scaling relational database service that scales up or down automatically based on the application demand. This service provides all the capabilities of Amazon Aurora, such as high availability, durability, and security, without requiring the customer to provision any database instances. With Amazon Aurora Serverless for MySQL, the sales team can enjoy minimal downtime since the database is designed to automatically scale to accommodate the increased traffic. Additionally, the service allows the customer to pay only for the capacity used, making it cost-effective for infrequent access patterns. Amazon RDS for MySQL could also be an option, but it requires the customer to select an instance type, and the database administrator would need to monitor and adjust the instance size manually to accommodate the increasing traffic.

NEW QUESTION 161
- (Topic 3)
A company needs to ingested and handle large amounts of streaming data that its application generates. The application runs on Amazon EC2 instances and sends data to Amazon Kinesis Data Streams. which is contained wild default settings. Every other day the application consumes the data and writes the data to an

Amazon S3 bucket for business intelligence (BI) processing the company observes that Amazon S3 is not receiving all the data that trio application sends to Kinesis Data Streams.
What should a solutions architect do to resolve this issue?

A. Update the Kinesis Data Streams default settings by modifying the data retention period.
B. Update the application to use the Kinesis Producer Library (KPL) lo send the data to Kinesis Data Streams.
C. Update the number of Kinesis shards lo handle the throughput of me data that is sent to Kinesis Data Streams.
D. Turn on S3 Versioning within the S3 bucket to preserve every version of every object that is ingested in the S3 bucket.

Answer: A

Explanation:
The data retention period of a Kinesis data stream is the time period from when a record is added to when it is no longer accessible1. The default retention period for a Kinesis data stream is 24 hours, which can be extended up to 8760 hours (365 days)1. The data retention period can be updated by using the AWS Management Console, the AWS CLI, or the Kinesis Data Streams API1.
To meet the requirements of the scenario, the solutions architect should update the Kinesis Data Streams default settings by modifying the data retention period. The solutions architect should increase the retention period to a value that is greater than or equal to the frequency of consuming the data and writing it to S32. This way, the company can ensure that S3 receives all the data that the application sends to Kinesis Data Streams.

NEW QUESTION 166
- (Topic 3)
A company collects data from a large number of participants who use wearabledevices.The company stores the data in an Amazon DynamoDB table and uses applications to analyze the data. The data workload is constant and predictable. The company wants to stay at or below its forecasted budget for DynamoDB.
Whihc solution will meet these requirements MOST cost-effectively?

A. Use provisioned mode and DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA). Reserve capacity for the forecasted workload.
B. Use provisioned mode Specify the read capacity units (RCUs) and write capacity units (WCUs).
C. Use on-demand mod
D. Set the read capacity unite (RCUs) and write capacity units (WCUs) high enough to accommodate changes in the workload.
E. Use on-demand mod
F. Specify the read capacity units (RCUs) and write capacity units (WCUs) with reserved capacity.

Answer: B

Explanation:
This option is the most efficient because it uses provisioned mode, which is a read/write capacity mode for processing reads and writes on your tables that lets you specify how much read and write throughput you expect your application to perform1. It also specifies the read capacity units (RCUs) and write capacity units (WCUs), which are the amount of data your application needs to read or write per second. It also meets the requirement of staying at or below its forecasted budget for DynamoDB, as provisioned mode has lower costs than on-demand mode for predictable workloads. This solution meets the requirement of collecting data from a large number of participants who use wearable devices with a constant and predictable data workload. Option A is less efficient because it uses provisioned mode and DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA), which is a storage class for infrequently accessed items that require milliseconds latency2. However, this does not meet the requirement of collecting data from a large number of participants who use wearable devices with a constant and predictable data workload, as DynamoDB Standard-IA is more suitable for items that are accessed less frequently than once every 30 days. Option C is less efficient because it uses on- demand mode, which is a read/write capacity mode that lets you pay only for what you use by automatically adjusting your table's capacity in response to changing demand3. However, this does not meet the requirement of staying at or below its forecasted budget for DynamoDB, as on-demand mode has higher costs than provisioned mode for predictable workloads. Option D is less efficient because it uses on-demand mode and specifies the RCUs and WCUs with reserved capacity, which is a way to reserve read and write capacity for your tables in exchange for discounted hourly rates. However, this does not meet the requirement of staying at or below its forecasted budget for DynamoDB, as on-demand mode has higher costs than provisioned mode for predictable workloads. Also, specifying RCUs and WCUs with reserved capacity is not possible with on-demand mode, as it only applies to provisioned mode.

NEW QUESTION 168
- (Topic 3)
An ecommerce company is building a distributed application that involves several serverless functions and AWS services to complete order-processing tasks. These tasks require manual approvals as part of the workflow A solutions architect needs to design an architecture for the order-processing application The solution must be able to combine multiple AWS Lambda functions into responsive serverless applications The solution also must orchestrate data and services that run on Amazon EC2 instances, containers, or on- premises servers
Which solution will meet these requirements with the LEAST operational overhead?

A. Use AWS Step Functions to build the application.
B. Integrate all the application components in an AWS Glue job
C. Use Amazon Simple Queue Service (Amazon SQS) to build the application
D. Use AWS Lambda functions and Amazon EventBridge (Amazon CloudWatch Events) events to build the application

Answer: A

Explanation:
AWS Step Functions is a fully managed service that makes it easy to build applications by coordinating the components of distributed applications and microservices using visual workflows. With Step Functions, you can combine multiple AWS Lambda functions into responsive serverless applications and orchestrate data and services that run on Amazon EC2 instances, containers, or on-premises servers. Step Functions also allows for manual approvals as part of the workflow. This solution meets all the requirements with the least operational overhead.
https://aws.amazon.com/step-functions/#:~:text=AWS%20Step%20Functions%20is%20a,machine%20learning%20(ML)%20pipelines.

NEW QUESTION 170
- (Topic 3)
A company stores confidential data in an Amazon Aurora PostgreSQL database in the ap- southeast-3 Region The database is encrypted with an AWS Key Management Service (AWS KMS) customer managed key The company was recently acquired and must securely share a backup of the database with the acquiring company's AWS account in ap- southeast-3.
What should a solutions architect do to meet these requirements?

A. Create a database snapshot Copy the snapshot to a new unencrypted snapshot Share the new snapshot with the acquiring company's AWS account
B. Create a database snapshot Add the acquiring company's AWS account to the KMS key policy Share the snapshot with the acquiring company's AWS account
C. Create a database snapshot that uses a different AWS managed KMS key Add the acquiring company's AWS account to the KMS key alia
D. Share the snapshot with the acquiring company's AWS account.
E. Create a database snapshot Download the database snapshot Upload the database snapshot to an Amazon S3 bucket Update the S3 bucket policy to allow access from the acquiring company's AWS account

Answer: B

Explanation:
https://docs.aws.amazon.com/kms/latest/developerguide/key-policy-modifying-external-accounts.html
There's no need to create another custom AWS KMS key. https://aws.amazon.com/premiumsupport/knowledge-center/aurora-share-encrypted- snapshot/ Give target account access to the custom AWS KMS key within the source account 1. Log in to the source account, and go to the AWS KMS console in the same Region as the DB cluster snapshot. 2. Select Customer-managed keys from the navigation pane. 3. Select your custom AWS KMS key (ALREADY CREATED) 4. From the Other AWS accounts section, select Add another AWS account, and then enter the AWS account number of your target account. Then: Copy and share the DB cluster snapshot

NEW QUESTION 175
- (Topic 3)
An application that is hosted on Amazon EC2 instances needs to access an Amazon S3 bucket Traffic must not traverse the internet How should a solutions architect configure access to meet these requirements?

A. Create a private hosted zone by using Amazon Route 53
B. Set up a gateway VPC endpoint for Amazon S3 in the VPC
C. Configure the EC2 instances to use a NAT gateway to access the S3 bucket
D. Establish an AWS Site-to-Site VPN connection between the VPC and the S3 bucket

Answer: B

Explanation:
This option is the most efficient because it uses a gateway VPC endpoint for Amazon S3, which provides reliable connectivity to Amazon S3 without requiring an internet gateway or a NAT device for the VPC1. A gateway VPC endpoint routes traffic from the VPC to Amazon S3 using a prefix list for the service and does not leave the AWS network2. This meets the requirement of not traversing the internet. Option A is less efficient because it uses a private hosted zone by using Amazon Route 53, which is a DNS service that allows you to create custom domain names for your resources within your VPC3. However, this does not provide connectivity to Amazon S3 without an internet gateway or a NAT device. Option C is less efficient because it uses a NAT gateway to access the S3 bucket, which is a highly available, managed Network Address Translation (NAT) service that enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating a connection with those instances4. However, this does not meet the requirement of not traversing the internet. Option D is less efficient because it uses an AWS Site-to-Site VPN connection between the VPC and the S3 bucket, which is a secure and encrypted network connection between your on-premises network and your VPC. However, this does not meet the requirement of not traversing the internet.

NEW QUESTION 180
- (Topic 3)
A financial company hosts a web application on AWS. The application uses an Amazon API Gateway Regional API endpoint to give users the ability to retrieve current stock prices. The company's security team has noticed an increase in the number of API requests. The security team is concerned that HTTP flood attacks might take the application offline.
A solutions architect must design a solution to protect the application from this type of attack.
Which solution meats these requirements with the LEAST operational overhead?

A. Create an Amazon CloudFront distribution in front of the API Gateway Regional API endpoint with a maximum TTL of 24 hours
B. Create a Regional AWS WAF web ACL with a rate-based rul
C. Associate the web ACL with the API Gateway stage.
D. Use Amazon CloudWatch metrics to monitor the Count metric and alert the security team when the predefined rate is reached
E. Create an Amazon CloudFront distribution with Lambda@Edge in front of the API Gateway Regional API endpoint Create an AWS Lambda function to block requests from IP addresses that exceed the predefined rate.

Answer: B

Explanation:
https://docs.aws.amazon.com/waf/latest/developerguide/web-acl.html
A rate-based rule in AWS WAF allows the security team to configure thresholds that trigger rate-based rules, which enable AWS WAF to track the rate of requests for a specified time period and then block them automatically when the threshold is exceeded. This provides the ability to prevent HTTP flood attacks with minimal operational overhead.

NEW QUESTION 183
- (Topic 3)
An ecommerce company is running a multi-tier application on AWS. The front-end and backend tiers run on Amazon EC2, and the database runs on Amazon RDS for MYSQL. The backend tier communities with the RDS instance. There are frequent calls to return identical database from the database that are causing performance slowdowns.
Which action should be taken to improve the performance of the backend?

A. Implement Amazon SNS to store the database calls.
B. Implement Amazon ElasticCache to cache the large database.
C. Implement an RDS for MySQL read replica to cache database calls.
D. Implement Amazon Kinesis Data Firehose to stream the calls to the database.

Answer: B

Explanation:
the best solution is to implement Amazon ElastiCache to cache the large datasets, which will store the frequently accessed data in memory, allowing for faster retrieval times. This can help to alleviate the frequent calls to the database, reduce latency, and improve the overall performance of the backend tier.

NEW QUESTION 188
- (Topic 3)
A company has implemented a self-managed DNS service on AWS. The solution consists of the following:
• Amazon EC2 instances in different AWS Regions
• Endpomts of a standard accelerator m AWS Global Accelerator
The company wants to protect the solution against DDoS attacks What should a solutions architect do to meet this requirement?

A. Subscribe to AWS Shield Advanced Add the accelerator as a resource to protect
B. Subscribe to AWS Shield Advanced Add the EC2 instances as resources to protect
C. Create an AWS WAF web ACL that includes a rate-based rule Associate the web ACL with the accelerator
D. Create an AWS WAF web ACL that includes a rate-based rule Associate the web ACL with the EC2 instances

Answer: A

Explanation:
AWS Shield is a managed service that provides protection against Distributed Denial of Service (DDoS) attacks for applications running on AWS. AWS Shield Standard is automatically enabled to all AWS customers at no additional cost. AWS Shield Advanced is an optional paid service. AWS Shield Advanced provides additional protections against more sophisticated and larger attacks for your applications running on Amazon Elastic Compute Cloud (EC2), Elastic Load Balancing (ELB), Amazon CloudFront, AWS Global Accelerator, and Route 53. https://docs.aws.amazon.com/waf/latest/developerguide/ddos-event-mitigation-logic- gax.html

NEW QUESTION 192
- (Topic 3)
A company runs an internal browser-based application The application runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. The Auto Scaling group scales up to 20 instances during work hours but scales down to 2 instances overnight Staff are complaining that the application is very slow when the day begins although it runs well by mid-morning.
How should the scaling be changed to address the staff complaints and keep costs to a minimum'?

A. Implement a scheduled action that sets the desired capacity to 20 shortly before the office opens
B. Implement a step scaling action triggered at a lower CPU threshold, and decrease the cooldown period.
C. Implement a target tracking action triggered at a lower CPU threshold, and decrease the cooldown period.
D. Implement a scheduled action that sets the minimum and maximum capacity to 20 shortly before the office opens

Answer: C

Explanation:
This option will scale up capacity faster in the morning to improve performance, but will still allow capacity to scale down during off hours. It achieves this as follows: • A target tracking action scales based on a CPU utilization target. By triggering at a lower CPU threshold in the morning, the Auto Scaling group will start scaling up sooner as traffic ramps up, launching instances before utilization gets too high and impacts performance. • Decreasing the cooldown period allows Auto Scaling to scale more aggressively, launching more instances faster until the target is reached. This speeds up the ramp-up of capacity. • However, unlike a scheduled action to set a fixed minimum/maximum capacity, with target tracking the group can still scale down during off hours based on demand. This helps minimize costs.

NEW QUESTION 195
- (Topic 3)
A company is using a fleet of Amazon EC2 instances to ingest data from on-premises data sources. The data is in JSON format and Ingestion rates can be as high as 1 MB/s. When an EC2 instance is rebooted, the data in-flight is lost. The company's data science team wants to query Ingested data In near-real time.
Which solution provides near-real -time data querying that is scalable with minimal data loss?

A. Publish data to Amazon Kinesis Data Streams Use Kinesis data Analytics to query the data.
B. Publish data to Amazon Kinesis Data Firehose with Amazon Redshift as the destination Use Amazon Redshift to query the data
C. Store ingested data m an EC2 Instance store Publish data to Amazon Kinesis Data Firehose with Amazon S3 as the destinatio
D. Use Amazon Athena to query the data.
E. Store ingested data m an Amazon Elastic Block Store (Amazon EBS) volume Publish data to Amazon ElastiCache tor Red Subscribe to the Redis channel to query the data

Answer: A

Explanation:
https://docs.aws.amazon.com/kinesisanalytics/latest/dev/what-is.html

NEW QUESTION 198
- (Topic 3)
A company runs a containerized application on a Kubernetes cluster in an on-premises data center. The company is using a MongoDB database for data storage. The company wants to migrate some of these environments to AWS, but no code changes or deployment method changes are possible at this time. The company needs a solution that minimizes operational overhead.
Which solution meets these requirements?

A. Use Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 worker nodes for compute and MongoDB on EC2 for data storage.
B. Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate for compute and Amazon DynamoDB for data storage.
C. Use Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 worker nodes for compute and Amazon DynamoDB for data storage.
D. Use Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate for compute and Amazon DocumentDB (with MongoDB compatibility) for data storage.

Answer: D

Explanation:
Amazon DocumentDB (with MongoDB compatibility) is a fast, reliable, and fully managed database service. Amazon DocumentDB makes it easy to set up, operate, and scale MongoDB-compatible databases in the cloud. With Amazon DocumentDB, you can run the same application code and use the same drivers and tools that you use with MongoDB. https://docs.aws.amazon.com/documentdb/latest/developerguide/what-is.html

NEW QUESTION 202
- (Topic 3)
A company's application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. On the first day of every month at midnight. The application becomes much slower when the month-end financial calcualtion bath runs. This causes the CPU utilization of the EC2 instaces to immediately peak to 100%, which disrupts the application.
What should a solution architect recommend to ensure the application is able to handle the workload and avoid downtime?

A. Configure an Amazon CloudFront distribution in from of the ALB.
B. Configure an EC2 Auto Scaling simple scaling policy based on CPU utilization.
C. Configure an EC2 Auto Scaling scheduled scaling policy based on the monthly schedule.
D. Configure Amazon ElasticCache to remove some of the workload from tha EC2 instances.

Answer: C

Explanation:
Configure an EC2 Auto Scaling scheduled scaling policy based on the monthly schedule is the best option because it allows for the proactive scaling of the EC2 instances before the monthly batch run begins. This will ensure that the application is able to handle the increased workload without experiencing downtime. The scheduled scaling policy can be configured to increase the number of instances in the Auto Scaling group a few hours before the batch run and then decrease the number of instances after the batch run is complete. This will ensure that the resources are available when needed and not wasted when not needed. The most appropriate solution to handle the increased workload during the monthly batch run and avoid downtime would be to configure an EC2 Auto Scaling scheduled scaling policy based on the monthly schedule. https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-scheduled- scaling.html

NEW QUESTION 204
......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons 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 FirstTry

* SAA-C03 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
Order The SAA-C03 Practice Test Here