Exam Questions AWS-Solution-Architect-Associate

Amazon AWS Certified Solutions Architect - Associate

https://www.2passeasy.com/dumps/AWS-Solution-Architect-Associate/

NEW QUESTION 1
- (Topic 4)
A company migrated a MySQL database from the company's on-premises data center to an Amazon RDS for MySQL DB instance. The company sized the RDS DB instance to meet the company's average daily workload. Once a month, the database performs slowly when the company runs queries for a report. The company wants to have the ability to run reports and maintain the performance of the daily workloads.
Which solution will meet these requirements?

A. Create a read replica of the databas
B. Direct the queries to the read replica.
C. Create a backup of the databas
D. Restore the backup to another DB instanc
E. Direct the queries to the new database.
F. Export the data to Amazon S3. Use Amazon Athena to query the S3 bucket.
G. Resize the DB instance to accommodate the additional workload.

Answer: C

Explanation:
Amazon Athena is a service that allows you to run SQL queries on data stored in Amazon S3. It is serverless, meaning you do not need to provision or manage any infrastructure. You only pay for the queries you run and the amount of data scanned1.
By using Amazon Athena to query your data in Amazon S3, you can achieve the following benefits:
? You can run queries for your report without affecting the performance of your
Amazon RDS for MySQL DB instance. You can export your data from your DB instance to an S3 bucket and use Athena to query the data in the bucket. This way, you can avoid the overhead and contention of running queries on your DB instance.
? You can reduce the cost and complexity of running queries for your report. You do
not need to create a read replica or a backup of your DB instance, which would incur additional charges and require maintenance. You also do not need to resize your DB instance to accommodate the additional workload, which would increase your operational overhead.
? You can leverage the scalability and flexibility of Amazon S3 and Athena. You can
store large amounts of data in S3 and query them with Athena without worrying about capacity or performance limitations. You can also use different formats, compression methods, and partitioning schemes to optimize your data storage and query performance1.

NEW QUESTION 2
- (Topic 4)
A company is creating an application that runs on containers in a VPC. The application stores and accesses data in an Amazon S3 bucket During the development phase, the application will store and access 1 TB of data in Amazon S3 each day. The company wants to minimize costs and wants to prevent traffic from traversing the internet whenever possible.
Which solution will meet these requirements?

A. Enable S3 Intelligent-Tiering for the S3 bucket.
B. Enable S3 Transfer Acceleration for the S3 bucket.
C. Create a gateway VPC endpoint for Amazon S3. Associate this endpoint with all route tables in the VPC.
D. Create an interface endpoint for Amazon S3 in the VP
E. Associate this endpoint with all route tables in the VPC.

Answer: C

Explanation:
A gateway VPC endpoint for Amazon S3 enables private connections between the VPC and Amazon S3 that do not require an internet gateway or NAT device. This minimizes costs and prevents traffic from traversing the internet. A gateway VPC endpoint uses a prefix list as the route target in a VPC route table to route traffic privately to Amazon S31. Associating the endpoint with all route tables in the VPC ensures that all subnets can access Amazon S3 through the endpoint.
Option A is incorrect because S3 Intelligent-Tiering is a storage class that optimizes storage costs by automatically moving objects between two access tiers based on changing access patterns. It does not affect the network traffic between the VPC and Amazon S32.
Option B is incorrect because S3 Transfer Acceleration is a feature that enables fast, easy, and secure transfers of files over long distances between clients and an S3 bucket. It does not prevent traffic from traversing the internet3.
Option D is incorrect because an interface VPC endpoint for Amazon S3 is powered by AWS PrivateLink, which requires an elastic network interface (ENI) with a private IP address in each subnet. This adds complexity and cost to the solution. Moreover, an interface VPC endpoint does not support cross-Region access to Amazon S3. Reference URL: 1: https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-s3.html 2: https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html#sc- dynamic-data-access 3: https://docs.aws.amazon.com/AmazonS3/latest/userguide/transfer-acceleration.html : https://aws.amazon.com/blogs/architecture/choosing-your-vpc-endpoint-strategy-for- amazon-s3/

NEW QUESTION 3
- (Topic 4)
A company has two VPCs named Management and Production. The Management VPC uses VPNs through a customer gateway to connect to a single device in the data center. The Production VPC uses a virtual private gateway AWS Direct Connect connections. The Management and Production VPCs both use a single VPC peering connection to allow communication between the
What should a solutions architect do to mitigate any single point of failure in this architecture?

A. Add a set of VPNs between the Management and Production VPCs.
B. Add a second virtual private gateway and attach it to the Management VPC.
C. Add a second set of VPNs to the Management VPC from a second customer gateway device.
D. Add a second VPC peering connection between the Management VPC and the Production VPC.

Answer: C

Explanation:
This answer is correct because it provides redundancy for the VPN connection between the Management VPC and the data center. If one customer gateway device or one VPN tunnel becomes unavailable, the traffic can still flow over the second customer gateway device and the second VPN tunnel. This way, the single point of failure in the VPN connection is mitigated.

References:
? https://docs.aws.amazon.com/vpn/latest/s2svpn/vpn-redundant-connection.html
? https://www.trendmicro.com/cloudoneconformity/knowledge-base/aws/VPC/vpn-tunnel-redundancy.html

NEW QUESTION 4
- (Topic 4)
A company wants to use an AWS CloudFormatlon stack for its application in a test environment. The company stores the CloudFormation template in an Amazon S3 bucket that blocks public access. The company wants to grant CloudFormation access to the template in the S3 bucket based on specific user requests to create the test environment The solution must follow security best practices.
Which solution will meet these requirements?

A. Create a gateway VPC endpoint for Amazon S3. Configure the CloudFormation stack to use the S3 object URL
B. Create an Amazon API Gateway REST API that has the S3 bucket as the targe
C. Configure the CloudFormat10n stack to use the API Gateway URL _
D. Create a presigned URL for the template object_ Configure the CloudFormation stack to use the presigned URL.
E. Allow public access to the template object in the S3 bucke
F. Block the public access after the test environment is created

Answer: C

Explanation:
it allows CloudFormation to access the template in the S3 bucket without granting public access or creating additional resources. A presigned URL is a URL that is signed with the access key of an IAM user or role that has permission to access the object. The presigned URL can be used by anyone who receives it, but it expires after a specified time. By creating a presigned URL for the template object and configuring the CloudFormation stack to use it, the company can grant CloudFormation access to the template based on specific user requests and follow security best practices. References:
? Using Amazon S3 Presigned URLs
? Using Amazon S3 Buckets

NEW QUESTION 5
- (Topic 4)
An ecommerce company is running a seasonal online sale. The company hosts its website on Amazon EC2 instances spanning multiple Availability Zones. The company wants its website to manage sudden traffic increases during the sale.
Which solution will meet these requirements MOST cost-effectively?

A. Create an Auto Scaling group that is large enough to handle peak traffic loa
B. Stop half of the Amazon EC2 instance
C. Configure the Auto Scaling group to use the stopped instances to scale out when traffic increases.
D. Create an Auto Scaling group for the websit
E. Set the minimum size of the Auto Scaling group so that it can handle high traffic volumes without the need to scale out.
F. Use Amazon CIoudFront and Amazon ElastiCache to cache dynamic content with an Auto Scaling group set as the origi
G. Configure the Auto Scaling group with the instances necessary to populate CIoudFront and ElastiCach
H. Scale in after the cache is fully populated.
I. Configure an Auto Scaling group to scale out as traffic increase
J. Create a launch template to start new instances from a preconfigured Amazon Machine Image (AMI).

Answer: D

Explanation:
The solution that meets the requirements of high availability, resiliency, and minimal operational effort is to use AWS Transfer for SFTP and an Amazon S3 bucket for storage. This solution allows the company to securely transfer files over SFTP to Amazon S3, which is a durable and scalable object storage service. The company can then modify the application to pull the batch files from Amazon S3 to an Amazon EC2 instance for processing. The EC2 instance can be part of an Auto Scaling group with a scheduled scaling policy to run the batch operation only at night. This way, the company can save costs by scaling down the EC2 instances when they are not needed. The other solutions do not meet all the requirements because they either use Amazon EFS or Amazon EBS for storage, which are more expensive and less scalable than Amazon S3, or they do not use a scheduled scaling policy to optimize the EC2 instances usage. References :=
? AWS Transfer for SFTP
? Amazon S3
? Amazon EC2 Auto Scaling

NEW QUESTION 6
- (Topic 4)
A company wants to analyze and generate reports to track the usage of its mobile app. The app is popular and has a global user base The company uses a custom report building program to analyze application usage.
The program generates multiple reports during the last week of each month. The program takes less than 10 minutes to produce each report. The company rarely uses the program to generate reports outside of the last week of each month. The company wants to generate reports in the least amount of time when the reports are requested.
Which solution will meet these requirements MOST cost-effectively?

A. Run the program by using Amazon EC2 On-Demand Instance
B. Create an Amazon EventBridge rule to start the EC2 instances when reports are requeste
C. Run the EC2 instances continuously during the last week of each month.
D. Run the program in AWS Lambd
E. Create an Amazon EventBridge rule to run a Lambda function when reports are requested.
F. Run the program in Amazon Elastic Container Service (Amazon ECS). Schedule Amazon ECS to run the program when reports are requested.
G. Run the program by using Amazon EC2 Spot Instance
H. Create an Amazon EventBridge rule to start the EC2 instances when reports are requeste
I. Run the EC2 instances continuously during the last week of each month.

Answer: B

Explanation:

This solution meets the requirements most cost-effectively because it leverages the serverless and event-driven capabilities of AWS Lambda and Amazon EventBridge. AWS Lambda allows you to run code without provisioning or managing servers, and you pay only for the compute time you consume. Amazon EventBridge is a serverless event bus service that lets you connect your applications with data from various sources and routes that data to targets such as AWS Lambda. By using Amazon EventBridge, you can create a rule that triggers a Lambda function to run the program when reports are requested, and you can also schedule the rule to run during the last week of each month. This way, you can generate reports in the least amount of time and pay only for the resources you use.
References:
? AWS Lambda
? Amazon EventBridge

NEW QUESTION 7
- (Topic 4)
A company needs to migrate a MySQL database from its on-premises data center to AWS within 2 weeks. The database is 20 TB in size. The company wants to complete the migration with minimal downtime.
Which solution will migrate the database MOST cost-effectively?

A. Order an AWS Snowball Edge Storage Optimized devic
B. Use AWS Database Migration Service (AWS DMS) with AWS Schema Conversion Tool (AWS SCT) to migrate the database with replication of ongoing change
C. Send the Snowball Edge device to AWS to finish the migration and continue the ongoing replication.
D. Order an AWS Snowmobile vehicl
E. Use AWS Database Migration Service (AWS DMS) with AWS Schema Conversion Tool (AWS SCT) to migrate the database wjgh ongoing change
F. Send the Snowmobile vehicle back to AWS to finish the migration and continue the ongoing replication.
G. Order an AWS Snowball Edge Compute Optimized with GPU devic
H. Use AWS Database Migration Service (AWS DMS) with AWS Schema Conversion Tool (AWS SCT) to migrate the database with ongoing change
I. Send the Snowball device to AWS to finish the migration and continue the ongoing replication.
J. Order a 1 GB dedicated AWS Direct Connect connection to establish a connection with the data cente
K. Use AWS Database Migration Service (AWS DMS) with AWS Schema Conversion Tool(AWS SCT) to migrate the database with replication of ongoing changes.

Answer: A

Explanation:
This answer is correct because it meets the requirements of migrating a 20 TB MySQL database within 2 weeks with minimal downtime and cost-effectively. The AWS Snowball Edge Storage Optimized device has up to 80 TB of usable storage space, which is enough to fit the database. The AWS Database Migration Service (AWS DMS) can migrate data from MySQL to Amazon Aurora, Amazon RDS for MySQL, or MySQL on Amazon EC2 with minimal downtime by continuously replicating changes from the source to the target. The AWS Schema Conversion Tool (AWS SCT) can convert the source schema and code to a format compatible with the target database. By using these services together, the company can migrate the database to AWS with minimal downtime and cost. The Snowball Edge device can be shipped back to AWS to finish the migration and continue the ongoing replication until the database is fully migrated.
References:
? https://docs.aws.amazon.com/snowball/latest/developer-guide/device- differences.html
? https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Source.MySQL.html
? https://docs.aws.amazon.com/SchemaConversionTool/latest/userguide/CHAP_So urce.MySQL.htm

NEW QUESTION 8
- (Topic 4)
A company's ecommerce website has unpredictable traffic and uses AWS Lambda functions to directly access a private Amazon RDS for PostgreSQL DB instance. The company wants to maintain predictable database performance and ensure that the Lambda invocations do not overload the database with too many connections.
What should a solutions architect do to meet these requirements?

A. Point the client driver at an RDS custom endpoin
B. Deploy the Lambda functions inside a VPC.
C. Point the client driver at an RDS proxy endpoin
D. Deploy the Lambda functions inside a VPC.
E. Point the client driver at an RDS custom endpoin
F. Deploy the Lambda functions outside a VPC.
G. Point the client driver at an RDS proxy endpoin
H. Deploy the Lambda functions outside a VPC.

Answer: B

Explanation:
To maintain predictable database performance and ensure that the Lambda invocations do not overload the database with too many connections, a solutions architect should point the client driver at an RDS proxy endpoint and deploy the Lambda functions inside a VPC. An RDS proxy is a fully managed database proxy that allows applications to share connections to a database, improving database availability and scalability. By using an RDS proxy, the Lambda functions can reuse existing connections, rather than creating new ones for every invocation, reducing the connection overhead and latency. Deploying the Lambda functions inside a VPC allows them to access the private RDS DB instance securely and efficiently, without exposing it to the public internet. References:
? Using Amazon RDS Proxy with AWS Lambda
? Configuring a Lambda function to access resources in a VPC

NEW QUESTION 9
- (Topic 4)
An image hosting company uploads its large assets to Amazon S3 Standard buckets The company uses multipart upload in parallel by using S3 APIs and overwrites if the same object is uploaded again. For the first 30 days after upload, the objects will be accessed frequently. The objects will be used less frequently after 30 days, but the access patterns for each object will be inconsistent The company must optimize its S3 storage costs while maintaining high availability and resiliency of stored assets.
Which combination of actions should a solutions architect recommend to meet these requirements? (Select TWO.)

A. Move assets to S3 Intelligent-Tiering after 30 days.
B. Configure an S3 Lifecycle policy to clean up incomplete multipart uploads.

C. Configure an S3 Lifecycle policy to clean up expired object delete markers.
D. Move assets to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days
E. Move assets to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 30 days.

Answer: AB

Explanation:
S3 Intelligent-Tiering is a storage class that automatically moves data to the most cost-effective access tier based on access frequency, without performance impact, retrieval fees, or operational overhead1. It is ideal for data with unknown or changing access patterns, such as the company's assets. By moving assets to S3 Intelligent-Tiering after 30 days, the company can optimize its storage costs while maintaining high availability and resilience of stored assets.
S3 Lifecycle is a feature that enables you to manage your objects so that they are stored cost effectively throughout their lifecycle2. You can create lifecycle rules to define actions that Amazon S3 applies to a group of objects. One of the actions is to abort incomplete multipart uploads that can occur when an upload is interrupted. By configuring an S3 Lifecycle policy to clean up incomplete multipart uploads, the company can reduce its storage costs and avoid paying for parts that are not used.
Option C is incorrect because expired object delete markers are automatically deleted by Amazon S3 and do not incur any storage costs3. Therefore, configuring an S3 Lifecycle policy to clean up expired object delete markers will not have any effect on the company's storage costs.
Option D is incorrect because S3 Standard-IA is a storage class for data that is accessed less frequently, but requires rapid access when needed1. It has a lower storage cost than S3 Standard, but it has a higher retrieval cost and a minimum storage duration charge of 30 days. Therefore, moving assets to S3 Standard-IA after 30 days may not optimize the company's storage costs if the assets are still accessed occasionally.
Option E is incorrect because S3 One Zone-IA is a storage class for data that is accessed less frequently, but requires rapid access when needed1. It has a lower storage cost than S3 Standard-IA, but it stores data in only one Availability Zone and has less resilience than other storage classes. It also has a higher retrieval cost and a minimum storage duration charge of 30 days. Therefore, moving assets to S3 One Zone-IA after 30 days may not optimize the company's storage costs if the assets are still accessed occasionally or require high availability. Reference URL: 1: https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html 2:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lifecycle-mgmt.html 3: https://docs.aws.amazon.com/AmazonS3/latest/userguide/delete-or-empty- bucket.html#delete-bucket-considerations : https://docs.aws.amazon.com/AmazonS3/latest/userguide/mpuoverview.html :
https://aws.amazon.com/certification/certified-solutions-architect-associate/

NEW QUESTION 10
- (Topic 4)
A company is creating an application The company stores data from tests of the application in multiple on-premises locations
The company needs to connect the on-premises locations to VPCs in an AWS Region in the AWS Cloud The number of accounts and VPCs will increase during the next year The network architecture must simplify the administration of new connections and must provide the ability to scale.
Which solution will meet these requirements with the LEAST administrative overhead'?

A. Create a peering connection between the VPCs Create a VPN connection between the VPCs and the on-premises locations.
B. Launch an Amazon EC2 instance On the instance, include VPN software that uses a VPN connection to connect all VPCs and on-premises locations.
C. Create a transit gateway Create VPC attachments for the VPC connections Create VPN attachments for the on-premises connections.
D. Create an AWS Direct Connect connection between the on-premises locations and acentral VP
E. Connect the central VPC to other VPCs by using peering connections.

Answer: C

Explanation:
A transit gateway is a network transit hub that enables you to connect your VPCs and on-premises networks in a centralized and scalable way. You can create VPC attachments to connect your VPCs to the transit gateway, and VPN attachments to connect your on-premises networks to the transit gateway over the internet. The transit gateway acts as a router between the attached networks, and simplifies the administration of new connections by reducing the number of peering or VPN connections required. You can also use transit gateway route tables to control the routing of traffic between the attached networks. By creating a transit gateway and using VPC and VPN attachments, you can meet the requirements of the company with the least administrative overhead.
References:
? AWS Transit Gateway
? Transit gateway attachments
? Transit gateway route tables

NEW QUESTION 10
- (Topic 4)
A company hosts a multi-tier web application on Amazon Linux Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The company observes that the Auto Scaling group launches more On-Demand Instances when the application's end users access high volumes of static web content. The company wants to optimize cost.
What should a solutions architect do to redesign the application MOST cost-effectively?

A. Update the Auto Scaling group to use Reserved Instances instead of On-Demand Instances.
B. Update the Auto Scaling group to scale by launching Spot Instances instead of On- Demand Instances.
C. Create an Amazon CloudFront distribution to host the static web contents from an Amazon S3 bucket.
D. Create an AWS Lambda function behind an Amazon API Gateway API to host the static website contents.

Answer: C

Explanation:
This answer is correct because it meets the requirements of optimizing cost and reducing the workload on the database. Amazon CloudFront is a content delivery network (CDN) service that speeds up distribution of your static and dynamic web content, such as .html,.css, .js, and image files, to your users. CloudFront delivers your content through a worldwide network of data centers called edge locations. When a user requests content that you're serving with CloudFront, the request is routed to the edge location that provides the lowest latency (time delay), so that content is delivered with the best possible performance. You can create an Amazon CloudFront distribution to host the static web contents from an Amazon S3 bucket, which is an origin that you define for CloudFront. This way, you can offload the requests for static web content from your EC2 instances to CloudFront, which can improve the performance and availability of your website, and reduce the cost of running your EC2 instances.
References:
? https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Introducti on.html
? https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html

NEW QUESTION 15
- (Topic 4)
A company wants to securely exchange data between its software as a service (SaaS) application Salesforce account and Amazon S3. The company must encrypt the data at rest by using AWS Key Management Service (AWS KMS) customer managed keys (CMKs). The company must also encrypt the data in transit. The company has enabled API access for the Salesforce account.
Which solution will meet these requirements with the LEAST development effort?

A. Create AWS Lambda functions to transfer the data securely from Salesforce to Amazon S3.
B. Create an AWS Step Functions workflow Define the task to transfer the data securely from Salesforce to Amazon S3.
C. Create Amazon AppFlow flows to transfer the data securely from Salesforce to Amazon S3.
D. Create a custom connector for Salesforce to transfer the data securely from Salesforce to Amazon S3.

Answer: C

Explanation:
Amazon AppFlow is a fully managed integration service that enables users to transfer data securely between SaaS applications and AWS services. It supports Salesforce as a source and Amazon S3 as a destination. It also supports encryption of data at rest using AWS KMS CMKs and encryption of data in transit using SSL/TLS1. By using Amazon AppFlow, the solution can meet the requirements with the least development effort.
* A. Create AWS Lambda functions to transfer the data securely from Salesforce to Amazon S3. This solution will not meet the requirement of the least development effort, as it involves writing custom code to interact with Salesforce and Amazon S3 APIs, handle authentication, encryption, error handling, and monitoring2.
* B. Create an AWS Step Functions workflow Define the task to transfer the data securely from Salesforce to Amazon S3. This solution will not meet the requirement of the least development effort, as it involves creating a state machine definition to orchestrate the data transfer task, and invoking Lambda functions or other services to perform the actual data transfer3.
* D. Create a custom connector for Salesforce to transfer the data securely from Salesforce to Ama-zon S3. This solution will not meet the requirement of the least development effort, as it involves using the Amazon AppFlow Custom Connector SDK to build and deploy a custom connector for Salesforce, which requires additional configuration and management. Reference URL: https://aws.amazon.com/appflow/

NEW QUESTION 20
- (Topic 4)
A company has an organization in AWS Organizations that has all features enabled The company requires that all API calls and logins in any existing or new AWS account must be audited The company needs a managed solution to prevent additional work and to minimize costs The company also needs to know when any AWS account is not compliant with the AWS Foundational Security Best Practices (FSBP) standard.
Which solution will meet these requirements with the LEAST operational overhead?

A. Deploy an AWS Control Tower environment in the Organizations management account Enable AWS Security Hub and AWS Control Tower Account Factory in the environment.
B. Deploy an AWS Control Tower environment in a dedicated Organizations member account Enable AWS Security Hub and AWS Control Tower Account Factory in the environment.
C. Use AWS Managed Services (AMS) Accelerate to build a multi-account landing zone (MALZ) Submit an RFC to self-service provision Amazon GuardDuty in the MALZ.
D. Use AWS Managed Services (AMS) Accelerate to build a multi-account landing zone (MALZ) Submit an RFC to self-service provision AWS Security Hub in the MALZ.

Answer: A

Explanation:
AWS Control Tower is a fully managed service that simplifies the setup and governance of a secure, compliant, multi-account AWS environment. It establishes a landing zone that is based on best-practices blueprints, and it enables governance using controls you can choose from a pre-packaged list. The landing zone is a well-architected, multi-account baseline that follows AWS best practices. Controls implement governance rules for security, compliance, and operations. AWS Security Hub is a service that provides a comprehensive view of your security posture across your AWS accounts. It aggregates, organizes, and prioritizes security alerts and findings from multiple AWS services, such as Amazon GuardDuty, Amazon Inspector, Amazon Macie, AWS Firewall Manager, and AWS IAM Access Analyzer, as well as from AWS Partner solutions. AWS Security Hub continuously monitors your environment using automated compliance checks based on the AWS best practices and industry standards, such as the AWS Foundational Security Best Practices (FSBP) standard. AWS Control Tower Account Factory is a feature that automates the provisioning of new AWS accounts that are preconfigured to meet your business, security, and compliance requirements. By deploying an AWS Control Tower environment in the Organizations management account, you can leverage the existing organization structure and policies, and enable AWS Security Hub and AWS Control Tower Account Factory in the environment. This way, you can audit all API calls and logins in any existing or new AWS account, monitor the compliance status of each account with the FSBP standard, and provision new accounts with ease and consistency. This solution meets the requirements with the least operational overhead, as you do not need to manage any infrastructure, perform any data migration, or submit any requests for changes. References:
? AWS Control Tower
? [AWS Security Hub]
? [AWS Control Tower Account Factory]

NEW QUESTION 25
- (Topic 4)
A company runs an application on AWS. The application receives inconsistent amounts of usage. The application uses AWS Direct Connect to connect to an on-premises MySQL- compatible database. The on-premises database consistently uses a minimum of 2 GiB of memory.
The company wants to migrate the on-premises database to a managed AWS service. The company wants to use auto scaling capabilities to manage unexpected workload increases.
Which solution will meet these requirements with the LEAST administrative overhead?

A. Provision an Amazon DynamoDB database with default read and write capacity settings.
B. Provision an Amazon Aurora database with a minimum capacity of 1 Aurora capacityunit (ACU).
C. Provision an Amazon Aurora Serverless v2 database with a minimum capacity of 1 Aurora capacity unit (ACU).
D. Provision an Amazon RDS for MySQL database with 2 GiB of memory.

Answer: C

Explanation:
it allows the company to migrate the on-premises database to a managed AWS service that supports auto scaling capabilities and has the least administrative

overhead. Amazon Aurora Serverless v2 is a configuration of Amazon Aurora that automatically scales compute capacity based on workload demand. It can scale from hundreds to hundreds of thousands of transactions in a fraction of a second. Amazon Aurora Serverless v2 also supports MySQL-compatible databases and AWS Direct Connect connectivity. References:
? Amazon Aurora Serverless v2
? Connecting to an Amazon Aurora DB Cluster

NEW QUESTION 26
- (Topic 4)
A company is deploying a new application to Amazon Elastic Kubernetes Service (Amazon EKS) with an AWS Fargate cluster. The application needs a storage solution for data persistence. The solution must be highly available and fault tolerant. The solution also must be shared between multiple application containers. Which solution will meet these requirements with the LEAST operational overhead?

A. Create Amazon Elastic Block Store (Amazon EBS) volumes in the same Availability Zones where EKS worker nodes are place
B. Register the volumes in a StorageClass object on an EKS cluste
C. Use EBS Multi-Attach to share the data between containers.
D. Create an Amazon Elastic File System (Amazon EFS) file syste
E. Register the file system in a StorageClass object on an EKS cluste
F. Use the same file system for all containers.
G. Create an Amazon Elastic Block Store (Amazon EBS) volum
H. Register the volume in a StorageClass object on an EKS cluste
I. Use the same volume for all containers.
J. Create Amazon Elastic File System (Amazon EFS) file systems in the same Availability Zones where EKS worker nodes are place
K. Register the file systems in a StorageClass object on an EKS cluste
L. Create an AWS Lambda function to synchronize the data between file systems.

Answer: B

Explanation:
Amazon EFS is a fully managed, elastic, and scalable file system that can be shared between multiple containers. It provides high availability and fault tolerance by
replicating data across multiple Availability Zones. Amazon EFS is compatible with Amazon EKS and AWS Fargate, and can be registered in a StorageClass object on an EKS cluster. Amazon EBS volumes are not supported by AWS Fargate, and cannot be shared between multiple containers without using EBS Multi-Attach, which has limitations and performance implications. EBS Multi-Attach also requires the volumes to be in the same Availability Zone as the worker nodes, which reduces availability and fault tolerance. Synchronizing data between multiple EFS file systems using AWS Lambda is unnecessary, complex, and prone to errors. References:
? Amazon EFS Storage Classes
? Amazon EKS Storage Classes
? Amazon EBS Multi-Attach

NEW QUESTION 27
- (Topic 4)
A company offers a food delivery service that is growing rapidly. Because of the growth, the company's order processing system is experiencing scaling problems during peak traffic hours. The current architecture includes the following:
• A group of Amazon EC2 instances that run in an Amazon EC2 Auto Scaling group to collect orders from the application
• Another group of EC2 instances that run in an Amazon EC2 Auto Scaling group to fulfill orders
The order collection process occurs quickly, but the order fulfillment process can take longer. Data must not be lost because of a scaling event.
A solutions architect must ensure that the order collection process and the order fulfillment process can both scale properly during peak traffic hours. The solution must optimize
utilization of the company's AWS resources. Which solution meets these requirements?

A. Use Amazon CloudWatch metrics to monitor the CPU of each instance in the Auto Scaling group
B. Configure each Auto Scaling group's minimum capacity according to peak workload values.
C. Use Amazon CloudWatch metrics to monitor the CPU of each instance in the Auto Scaling group
D. Configure a CloudWatch alarm to invoke an Amazon Simple Notification Service (Amazon SNS) topic that creates additional Auto Scaling groups on demand.
E. Provision two Amazon Simple Queue Service (Amazon SQS) queues: one for order collection and another for order fulfillmen
F. Configure the EC2 instances to poll their respective queu
G. Scale the Auto Scaling groups based on notifications that the queues send.
H. Provision two Amazon Simple Queue Service (Amazon SQS) queues: one for order collection and another for order fulfillmen
I. Configure the EC2 instances to poll their respective queu
J. Create a metric based on a backlog per instance calculatio
K. Scale the Auto Scaling groups based on this metric.

Answer: D

Explanation:
The number of instances in your Auto Scaling group can be driven by how long it takes to process a message and the acceptable amount of latency (queue delay). The solution is to use a backlog per instance metric with the target value being the acceptable backlog per instance to maintain.

NEW QUESTION 31
- (Topic 4)
A company has established a new AWS account. The account is newly provisioned and no changes have been made to the default settings. The company is concerned about the security of the AWS account root user.
What should be done to secure the root user?

A. Create 1AM users for daily administrative task
B. Disable the root user.
C. Create 1AM users for daily administrative task
D. Enable multi-factor authentication on the root user.
E. Generate an access key for the root user Use the access key for daily administration tasks instead of the AWS Management Console.
F. Provide the root user credentials to the most senior solutions architec

G. Have the solutions architect use the root user for daily administration tasks.

Answer: B

Explanation:
This answer is the most secure and recommended option for securing the root user of a new AWS account. The root user is the identity that has complete access to all AWS services and resources in the account. It is accessed by signing in with the email address and password that were used to create the account. To protect the root user credentials from unauthorized use, AWS advises the following best practices:
? Create IAM users for daily administrative tasks. IAM users are identities that you
create in your account that have specific permissions to access AWS resources. You can create individual IAM users for yourself and for others who need access to your account. You can also assign IAM users to IAM groups that have a set of policies that grant permissions to perform common tasks. By using IAM users instead of the root user, you can follow the principle of least privilege and reduce the risk of compromising your account.
? Enable multi-factor authentication (MFA) on the root user. MFA is a security
feature that requires users to prove their identity by providing two pieces of information: their password and a code from a device that only they have access to. By enabling MFA on the root user, you can add an extra layer of protection to your account and prevent unauthorized access even if your password is compromised.
? Limit the tasks you perform with the root user account. You should use the root
user only for tasks that require root user credentials, such as changing your account settings, closing your account, or managing consolidated billing. For a complete list of tasks that require root user credentials, see Tasks that require root user credentials. For all other tasks, you should use IAM users or roles that have the appropriate permissions.
References:
? AWS account root user
? Root user best practices for your AWS account
? Tasks that require root user credentials

NEW QUESTION 32
- (Topic 4)
A company is using an Application Load Balancer (ALB) to present its application to the internet. The company finds abnormal traffic access patterns across the application. A solutions architect needs to improve visibility into the infrastructure to help the company understand these abnormalities better.
What is the MOST operationally efficient solution that meets these requirements?

A. Create a table in Amazon Athena for AWS CloudTrail log
B. Create a query for the relevant information.
C. Enable ALB access logging to Amazon S3. Create a table in Amazon Athena, and query the logs.
D. Enable ALB access logging to Amazon S3 Open each file in a text editor, and search each line for the relevant information
E. Use Amazon EMR on a dedicated Amazon EC2 instance to directly query the ALB to acquire traffic access log information.

Answer: B

Explanation:
This solution meets the requirements because it allows the company to improve visibility into the infrastructure by using ALB access logging and Amazon Athena. ALB access logging is a feature that captures detailed information about requests sent to the load balancer, such as the client's IP address, request path, response code, and latency. By enabling ALB access logging to Amazon S3, the company can store the access logs in an S3 bucket as compressed files. Amazon Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL. By creating a table in Amazon Athena for the access logs, the company can query the logs and get results in seconds. This way, the company can better understand the abnormal traffic access patterns across the application.
References:
? Access logs for your Application Load Balancer
? Querying Application Load Balancer Logs

NEW QUESTION 34
- (Topic 4)
A company wants to migrate its three-tier application from on premises to AWS. The web tier and the application tier are running on third-party virtual machines (VMs). The database tier is running on MySQL.
The company needs to migrate the application by making the fewest possible changes to the architecture. The company also needs a database solution that can restore data to a specific point in time.
Which solution will meet these requirements with the LEAST operational overhead?

A. Migrate the web tier and the application tier to Amazon EC2 instances in private subnet
B. Migrate the database tier to Amazon RDS for MySQL in private subnets.
C. Migrate the web tier to Amazon EC2 instances in public subnet
D. Migrate the application tier to EC2 instances in private subnet
E. Migrate the database tier to Amazon Aurora MySQL in private subnets.
F. Migrate the web tier to Amazon EC2 instances in public subnet
G. Migrate the application tier to EC2 instances in private subnet
H. Migrate the database tier to Amazon RDS for MySQL in private subnets.
I. Migrate the web tier and the application tier to Amazon EC2 instances in public subnet
J. Migrate the database tier to Amazon Aurora MySQL in public subnets.

Answer: C

Explanation:
The solution that meets the requirements with the least operational overhead is to migrate the web tier to Amazon EC2 instances in public subnets, migrate the application tier to EC2 instances in private subnets, and migrate the database tier to Amazon RDS for MySQL in private subnets. This solution allows the company to migrate its three-tier application to AWS by making minimal changes to the architecture, as it preserves the same web, application, and database tiers and uses the same MySQL database engine. The solution also provides a database solution that can restore data to a specific point in time, as Amazon RDS for MySQL supports automated backups and point- in-time recovery. This solution also reduces the operational overhead by using managed services such as Amazon EC2 and Amazon RDS, which handle tasks such as provisioning, patching, scaling, and monitoring.
The other solutions do not meet the requirements as well as the first one because they either involve more changes to the architecture, do not provide point-in-time recovery, or do not follow best practices for security and availability. Migrating the database tier to Amazon Aurora MySQL would require changing the database engine and potentially modifying the application code to ensure compatibility. Migrating the web tier and the application tier to public subnets would expose them to more security risks and reduce their availability in case of a subnet failure. Migrating the database tier to public subnets would also compromise its security and

performance. References:
? Migrate Your Application Database to Amazon RDS
? Amazon RDS for MySQL
? Amazon Aurora MySQL
? Amazon VPC

NEW QUESTION 37
- (Topic 4)
A company used an Amazon RDS for MySQL DB instance during application testing. Before terminating the DB instance at the end of the test cycle, a solutions architect created two backups. The solutions architect created the first backup by using the mysqldump utility to create a database dump. The solutions architect created the second backup by enabling the final DB snapshot option on RDS termination.
The company is now planning for a new test cycle and wants to create a new DB instance from the most recent backup. The company has chosen a MySQL-compatible edition of Amazon Aurora to host the DB instance.
Which solutions will create the new DB instance? (Select TWO.)

A. Import the RDS snapshot directly into Aurora.
B. Upload the RDS snapshot to Amazon S3. Then import the RDS snapshot into Aurora.
C. Upload the database dump to Amazon S3. Then import the database dump into Aurora.
D. Use AWS Database Migration Service (AWS DMS) to import the RDS snapshot into Aurora.
E. Upload the database dump to Amazon S3. Then use AWS Database Migration Service (AWS DMS) to import the database dump into Aurora.

Answer: AC

Explanation:
These answers are correct because they meet the requirements of creating a new DB instance from the most recent backup and using a MySQL-compatible edition of Amazon Aurora to host the DB instance. You can import the RDS snapshot directly into Aurora if the MySQL DB instance and the Aurora DB cluster are running the same version of MySQL. For example, you can restore a MySQL version 5.6 snapshot directly to Aurora MySQL version 5.6, but you can't restore a MySQL version 5.6 snapshot directly to Aurora MySQL version 5.7. This method is simple and requires the fewest number of steps. You can upload the database dump to Amazon S3 and then import the database dump into Aurora if the MySQL DB instance and the Aurora DB cluster are running different versions of MySQL. For example, you can import a MySQL version 5.6 database dump into Aurora MySQL version 5.7, but you can't restore a MySQL version 5.6 snapshot directly to Aurora MySQL version 5.7. This method is more flexible and allows you to migrate across different versions of MySQL.
References:
? https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL
.Migrating.RDSMySQL.Import.html
? https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL
.Migrating.RDSMySQL.Dump.html

NEW QUESTION 39
- (Topic 4)
A company runs a web application that is deployed on Amazon EC2 instances in the private subnet of a VPC. An Application Load Balancer (ALB) that extends across the public subnets directs web traffic to the EC2 instances. The company wants to implement new security measures to restrict inbound traffic from the ALB to the EC2 instances while preventing access from any other source inside or outside the private subnet of the EC2 instances.
Which solution will meet these requirements?

A. Configure a route in a route table to direct traffic from the internet to the private IP addresses of the EC2 instances.
B. Configure the security group for the EC2 instances to only allow traffic that comes from the security group for the ALB.
C. Move the EC2 instances into the public subne
D. Give the EC2 instances a set of Elastic IP addresses.
E. Configure the security group for the ALB to allow any TCP traffic on any port.

Answer: B

Explanation:
To restrict inbound traffic from the ALB to the EC2 instances, the security group for the EC2 instances should only allow traffic that comes from the security group for the ALB. This way, the EC2 instances can only receive requests from the ALB and not from any other source inside or outside the private subnet.
References:
? Security Groups for Your Application Load Balancers
? Security Groups for Your VPC

NEW QUESTION 43
- (Topic 4)
A company has deployed a multiplayer game for mobile devices. The game requires live
location tracking of players based on latitude and longitude. The data store for the game must support rapid updates and retrieval of locations.
The game uses an Amazon RDS for PostgreSQL DB instance with read replicas to store the location data. During peak usage periods, the database is unable to maintain the performance that is needed for reading and writing updates. The game's user base is increasing rapidly.
What should a solutions architect do to improve the performance of the data tier?

A. Take a snapshot of the existing DB instanc
B. Restore the snapshot with Multi-AZ enabled.
C. Migrate from Amazon RDS to Amazon OpenSearch Service with OpenSearch Dashboards.
D. Deploy Amazon DynamoDB Accelerator (DAX) in front of the existing DB instanc
E. Modify the game to use DAX.
F. Deploy an Amazon ElastiCache for Redis cluster in front of the existing DB instanc
G. Modify the game to use Redis.

Answer: D

Explanation:
The solution that will improve the performance of the data tier is to deploy an Amazon ElastiCache for Redis cluster in front of the existing DB instance and modify the game to use Redis. This solution will enable the game to store and retrieve the location data of the players in a fast and scalable way, as Redis is an in-

memory data store that supports geospatial data types and commands. By using ElastiCache for Redis, the game can reduce the load on the RDS for PostgreSQL DB instance, which is not optimized for high-frequency updates and queries of location data. ElastiCache for Redis also supports replication, sharding, and auto scaling to handle the increasing user base of the game. The other solutions are not as effective as the first one because they either do not improve the performance, do not support geospatial data, or do not leverage caching. Taking a snapshot of the existing DB instance and restoring it with Multi-AZ enabled will not improve the performance of the data tier, as it only provides high availability and durability, but not scalability or low latency. Migrating from Amazon RDS to Amazon OpenSearch Service with OpenSearch Dashboards will not improve the performance of the data tier, as OpenSearch Service is mainly designed for full-text search and analytics, not for real-time location tracking. OpenSearch Service also does not support geospatial data types and commands natively, unlike Redis. Deploying Amazon DynamoDB Accelerator (DAX) in front of the existing DB instance and modifying the game to use DAX will not improve the performance of the data tier, as DAX is only compatible with DynamoDB, not with RDS for PostgreSQL. DAX also does not support geospatial data types and commands.
References:
? Amazon ElastiCache for Redis
? Geospatial Data Support - Amazon ElastiCache for Redis
? Amazon RDS for PostgreSQL
? Amazon OpenSearch Service
? Amazon DynamoDB Accelerator (DAX)

NEW QUESTION 44
- (Topic 4)
A company is conducting an internal audit. The company wants to ensure that the data in an Amazon S3 bucket that is associated with the company's AWS Lake Formation data lake does not contain sensitive customer or employee data. The company wants to discover personally identifiable information (Pll) or financial information, including passport numbers and credit card numbers.
Which solution will meet these requirements?

A. Configure AWS Audit Manager on the accoun
B. Select the Payment Card Industry Data Security Standards (PCI DSS) for auditing.
C. Configure Amazon S3 Inventory on the S3 bucke
D. Configure Amazon Athena to query the inventory.
E. Configure Amazon Macie to run a data discovery job that uses managed identifiers for the required data types.
F. Use Amazon S3 Select to run a report across the S3 bucket.

Answer: C

Explanation:
Amazon Macie is a fully managed data security and data privacy service that uses machine learning and pattern matching to discover and protect your sensitive data in AWS. Macie can run data discovery jobs that use managed identifiers for various types of PII or financial information, such as passport numbers and credit card numbers. Macie can also generate findings that alert you to potential issues or risks with your data. References:
https://docs.aws.amazon.com/macie/latest/userguide/macie-identifiers.html

NEW QUESTION 47
- (Topic 4)
A company wants to host a scalable web application on AWS. The application will be accessed by users from different geographic regions of the world. Application users will be able to download and upload unique data up to gigabytes in size. The development team wants a cost-effective solution to minimize upload and download latency and maximize performance.
What should a solutions architect do to accomplish this?

A. Use Amazon S3 with Transfer Acceleration to host the application.
B. Use Amazon S3 with CacheControl headers to host the application.
C. Use Amazon EC2 with Auto Scaling and Amazon CloudFront to host the application.
D. Use Amazon EC2 with Auto Scaling and Amazon ElastiCache to host the application.

Answer: C

Explanation:
This answer is correct because it meets the requirements of hosting a scalable web application that can handle large data transfers from different geographic regions. Amazon EC2 provides scalable compute capacity for hosting web applications. Auto Scaling can automatically adjust the number of EC2 instances based on the demand and traffic patterns. Amazon CloudFront is a content delivery network (CDN) that can cache static and dynamic content at edge locations closer to the users, reducing latency and improving performance. CloudFront can also use S3 Transfer Acceleration to speed up the transfers between S3 buckets and CloudFront edge locations.
References:
? https://docs.aws.amazon.com/autoscaling/ec2/userguide/what-is-amazon-ec2- auto-scaling.html
? https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Introducti on.html
? https://aws.amazon.com/s3/transfer-acceleration/

NEW QUESTION 49
- (Topic 4)
A solutions architect has created two IAM policies: Policy1 and Policy2. Both policies are attached to an IAM group.

Policy 1

```
{
  "Version": "2012-10-17",  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iam:Get*",
        "iam:List*",
        "kms:List*",
        "ec2:*",
        "ds:*",
        "logs:Get*",
        "logs:Describe*"
      ],
      "Resource": "*"
    }
  ]
}
```

Policy 2

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "ds:Delete*",
      "Resource": "*"
    }
  ]
}
```

A cloud engineer is added as an IAM user to the IAM group. Which action will the cloud engineer be able to perform?

A. Deleting IAM users
B. Deleting directories
C. Deleting Amazon EC2 instances
D. Deleting logs from Amazon CloudWatch Logs

Answer: C

Explanation:
https://awscli.amazonaws.com/v2/documentation/api/latest/reference/ds/index.html

NEW QUESTION 54
- (Topic 4)
A retail company has several businesses. The IT team for each business manages its own AWS account. Each team account is part of an organization in AWS Organizations. Each team monitors its product inventory levels in an Amazon DynamoDB table in the team's own AWS account.
The company is deploying a central inventory reporting application into a shared AWS account. The application must be able to read items from all the teams' DynamoDB tables.
Which authentication option will meet these requirements MOST securely?

A. Integrate DynamoDB with AWS Secrets Manager in the inventory application accoun
B. Configure the application to use the correct secret from Secrets Manager to authenticate and read the DynamoDB tabl
C. Schedule secret rotation for every 30 days.
D. In every business account, create an 1AM user that has programmatic acces
E. Configure the application to use the correct 1AM user access key ID and secret access key to authenticate and read the DynamoDB tabl
F. Manually rotate 1AM access keys every 30 days.
G. In every business account, create an 1AM role named BU_ROLE with a policy that gives the role access to the DynamoDB table and a trust policy to trust a specific role in the inventory application accoun
H. In the inventory account, create a role named APP_ROLE that allows access to the STS AssumeRole API operatio
I. Configure the application to use APP_ROLE and assume the cross-account role BU_ROLE to read the DynamoDB table.
J. Integrate DynamoDB with AWS Certificate Manager (ACM). Generate identity certificates to authenticate DynamoD
K. Configure the application to use the correct certificate to authenticate and read the DynamoDB table.

Answer: C

Explanation:
This solution meets the requirements most securely because it uses IAM roles and the STS AssumeRole API operation to authenticate and authorize the inventory application to access the DynamoDB tables in different accounts. IAM roles are more secure than IAM users or certificates because they do not require long-term credentials or passwords. Instead, IAM roles provide temporary security credentials that are automatically rotated and can be configured with a limited duration. The STS AssumeRole API operation enables you to request temporary credentials for a role that you are allowed to assume. By using this operation, you

can delegate access to resources that are in different AWS accounts that you own or that are owned by third parties. The trust policy of the role defines which entities can assume the role, and the permissions policy of the role defines which actions can be performed on the resources. By using this solution, you can avoid hard- coding credentials or certificates in the inventory application, and you can also avoid storing them in Secrets Manager or ACM. You can also leverage the built-in security features of IAM and STS, such as MFA, access logging, and policy conditions.
References:
? IAM Roles
? STS AssumeRole
? Tutorial: Delegate Access Across AWS Accounts Using IAM Roles

NEW QUESTION 59
- (Topic 4)
A company is making a prototype of the infrastructure for its new website by manually provisioning the necessary infrastructure. This infrastructure includes an Auto Scaling group, an Application Load Balancer, and an Amazon RDS database. After the configuration has been thoroughly validated, the company wants the capability to immediately deploy the infrastructure for development and production use in two
Availability Zones in an automated fashion.
What should a solutions architect recommend to meet these requirements?

A. Use AWS Systems Manager to replicate and provision the prototype infrastructure in two Availability Zones.
B. Define the infrastructure as a template by using the prototype infrastructure as a guid
C. Deploy the infrastructure with AWS CloudFormation
D. Use AWS Config to record the inventory of resources that are used in the prototype infrastructur
E. Use AWS Config to deploy the prototype infrastructure into two Availability Zones.
F. Use AWS Elastic Beanstalk and configure it to use an automated reference to the prototype infrastructure to automatically deploy new environments in two Availability Zones

Answer: B

Explanation:
AWS CloudFormation is a service that helps you model and set up your AWS resources by using templates that describe all the resources that you want, such as Auto Scaling groups, load balancers, and databases. You can use AWS CloudFormation to deploy your infrastructure in an automated and consistent way across multiple environments and regions. You can also use AWS CloudFormation to update or delete your infrastructure as a single unit.
Reference URLs:
1 https://aws.amazon.com/cloudformation/
2 https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/Welcome.html
3 https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-whatis- concepts.html

NEW QUESTION 62
- (Topic 4)
A company runs a three-tier application in two AWS Regions. The web tier, the application tier, and the database tier run on Amazon EC2 instances. The company uses Amazon RDS for Microsoft SQL Server Enterprise for the database tier The database tier is experiencing high load when weekly and monthly reports are run. The company wants to reduce the load on the database tier.
Which solution will meet these requirements with the LEAST administrative effort?

A. Create read replica
B. Configure the reports to use the new read replicas.
C. Convert the RDS database to Amazon DynamoDB_ Configure the reports to use DynamoDB
D. Modify the existing RDS DB instances by selecting a larger instance size.
E. Modify the existing ROS DB instances and put the instances into an Auto Scaling group.

Answer: A

Explanation:
it allows the company to create read replicas of its RDS database and reduce the load on the database tier. By creating read replicas, the company can offload read traffic from the primary database instance to one or more replicas. By configuring the reports to use the new read replicas, the company can improve performance and availability of its database tier. References:
? Working with Read Replicas
? Read Replicas for Amazon RDS for SQL Server

NEW QUESTION 66
- (Topic 4)
A company hosts a data lake on Amazon S3. The data lake ingests data in Apache Parquet format from various data sources. The company uses multiple transformation steps to prepare the ingested data. The steps include filtering of anomalies, normalizing of data to standard date and time values, and generation of aggregates for analyses.
The company must store the transformed data in S3 buckets that data analysts access. The company needs a prebuilt solution for data transformation that does not require code. The solution must provide data lineage and data profiling. The company needs to share the data transformation steps with employees throughout the company.
Which solution will meet these requirements?

A. Configure an AWS Glue Studio visual canvas to transform the dat
B. Share the transformation steps with employees by using AWS Glue jobs.
C. Configure Amazon EMR Serverless to transform the dat
D. Share the transformation steps with employees by using EMR Serveriess jobs.
E. Configure AWS Glue DataBrew to transform the dat
F. Share the transformation steps with employees by using DataBrew recipes.
G. Create Amazon Athena tables for the dat
H. Write Athena SQL queries to transform the dat
I. Share the Athena SQL queries with employees.

Answer: C

Explanation:
The most suitable solution for the company's requirements is to configure AWS Glue DataBrew to transform the data and share the transformation steps with employees by using DataBrew recipes. This solution will provide a prebuilt solution for data transformation that does not require code, and will also provide data lineage and data profiling. The company can easily share the data transformation steps with employees throughout the company by using DataBrew recipes.
AWS Glue DataBrew is a visual data preparation tool that makes it easy for data analysts and data scientists to clean and normalize data for analytics or machine learning by up to 80% faster. Users can upload their data from various sources, such as Amazon S3, Amazon RDS, Amazon Redshift, Amazon Aurora, or Glue Data Catalog, and use a point- and-click interface to apply over 250 built-in transformations. Users can also preview the results of each transformation step and see how it affects the quality and distribution of the data1.
A DataBrew recipe is a reusable set of transformation steps that can be applied to one or more datasets. Users can create recipes from scratch or use existing ones from the DataBrew recipe library. Users can also export, import, or share recipes with other users or groups within their AWS account or organization2.
DataBrew also provides data lineage and data profiling features that help users understand and improve their data quality. Data lineage shows the source and destination of each dataset and how it is transformed by each recipe step. Data profiling shows various statistics and metrics about each dataset, such as column

NEW QUESTION 70
- (Topic 4)
A company operates an ecommerce website on Amazon EC2 instances behind an Application Load Balancer (ALB) in an Auto Scaling group. The site is experiencing performance issues related to a high request rate from illegitimate external systems with changing IP addresses. The security team is worried about potential DDoS attacks against the website. The company must block the illegitimate incoming requests in a way that has a minimal impact on legitimate users.
What should a solutions architect recommend?

A. Deploy Amazon Inspector and associate it with the ALB.
B. Deploy AWS WAF, associate it with the ALB, and configure a rate-limiting rule.
C. Deploy rules to the network ACLs associated with the ALB to block the incoming traffic.
D. Deploy Amazon GuardDuty and enable rate-limiting protection when configuring GuardDuty.

Answer: B

Explanation:
This answer is correct because it meets the requirements of blocking the illegitimate incoming requests in a way that has a minimal impact on legitimate users. AWS WAF is a web application firewall that helps protect your web applications or APIs against common web exploits that may affect availability, compromise security, or consume excessive resources. AWS WAF gives you control over how traffic reaches your applications by enabling you to create security rules that block common attack patterns, such as SQL injection or cross-site scripting, and rules that filter out specific traffic patterns you define. You can associate AWS WAF with an ALB to protect the web application from malicious requests. You can configure a rate-limiting rule in AWS WAF to track the rate of requests for each originating IP address and block requests from an IP address that exceeds a certain limit within a five-minute period. This way, you can mitigate potential DDoS attacks and improve the performance of your website.
References:
? https://docs.aws.amazon.com/waf/latest/developerguide/what-is-aws-waf.html
? https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type- rate-based.html

NEW QUESTION 74
- (Topic 4)
A company designed a stateless two-tier application that uses Amazon EC2 in a single Availability Zone and an Amazon RDS Multi-AZ DB instance New company management wants to ensure the application is highly available.
What should a solutions architect do to meet this requirement?

A. Configure the application to use Multi-AZ EC2 Auto Scaling and create an Application Load Balancer
B. Configure the application to take snapshots of the EC2 instances and send them to a different AWS Region.
C. Configure the application to use Amazon Route 53 latency-based routing to feed requests to the application.
D. Configure Amazon Route 53 rules to handle incoming requests and create a Multi-AZ Application Load Balancer

Answer: A

Explanation:
https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-add-availability- zone.html

NEW QUESTION 76
- (Topic 4)
A gaming company wants to launch a new internet-facing application in multiple AWS Regions The application will use the TCP and UDP protocols for communication. The company needs to provide high availability and minimum latency for global users.
Which combination of actions should a solutions architect take to meet these requirements? (Select TWO.)

A. Create internal Network Load Balancers in front of the application in each Region.
B. Create external Application Load Balancers in front of the application in each Region.
C. Create an AWS Global Accelerator accelerator to route traffic to the load balancers in each Region.
D. Configure Amazon Route 53 to use a geolocation routing policy to distribute the traffic.
E. Configure Amazon CloudFront to handle the traffic and route requests to the application in each Region.

Answer: BC

Explanation:
This combination of actions will provide high availability and minimum latency for global users by using AWS Global Accelerator and Application Load Balancers. AWS Global Accelerator is a networking service that helps you improve the availability, performance, and security of your internet-facing applications by using the AWS global network. It provides two global static public IPs that act as a fixed entry point to your application endpoints, such as Application Load Balancers, in multiple Regions1. Global Accelerator uses the AWS backbone network to route traffic to the optimal regional
endpoint based on health, client location, and policies that you configure. It also offers TCP and UDP support, traffic encryption, and DDoS protection2. Application Load Balancers are external load balancers that distribute incoming application traffic across multiple targets, such as EC2 instances, in multiple Availability Zones. They support both HTTP and HTTPS (SSL/TLS) protocols, and offer advanced features such as content-based routing, health checks, and integration with other AWS services3. By creating external Application Load Balancers in front of the application in each Region, you can ensure that the application can handle varying load patterns and scale on demand. By creating an AWS Global Accelerator accelerator to route traffic to the load balancers in each Region, you can leverage the performance, security, and availability of the AWS global network to deliver the best possible user experience.

References: 1: What is AWS Global Accelerator? - AWS Global Accelerator4, Overview section2: Network Acceleration Service - AWS Global Accelerator - AWS5, Why AWS Global Accelerator? section. 3: What is an Application Load Balancer? - Elastic Load Balancing6, Overview section.

NEW QUESTION 77
- (Topic 4)
A company uses Amazon S3 as its data lake. The company has a new partner that must use SFTP to upload data files A solutions architect needs to implement a highly available SFTP solution that minimizes operational overhead.
Which solution will meet these requirements?

A. Use AWS Transfer Family to configure an SFTP-enabled server with a publicly accessible endpoint Choose the S3 data lake as the destination
B. Use Amazon S3 File Gateway as an SFTP server Expose the S3 File Gateway endpoint URL to the new partner Share the S3 File Gateway endpoint with the newpartner
C. Launch an Amazon EC2 instance in a private subnet in a VP
D. Instruct the new partner to upload files to the EC2 instance by using a VP
E. Run a cron job script on the EC2 instance to upload files to the S3 data lake
F. Launch Amazon EC2 instances in a private subnet in a VP
G. Place a Network Load Balancer (NLB) in front of the EC2 instance
H. Create an SFTP listener port for the NLBShare the NLB hostname with the new partner Run a cron job script on the EC2 instances to upload files to the S3 data lake.

Answer: A

Explanation:
This option is the most cost-effective and simple way to enable SFTP access to the S3 data lake. AWS Transfer Family is a fully managed service that supports secure file transfers over SFTP, FTPS, and FTP protocols. You can create an SFTP-enabled server with a public endpoint and associate it with your S3 bucket. You can also use AWS Identity and Access Management (IAM) roles and policies to control access to your S3 data lake. The service scales automatically to handle any volume of file transfers and provides high availability and durability. You do not need to provision, manage, or patch any servers or load balancers.
Option B is not correct because Amazon S3 File Gateway is not an SFTP server. It is a hybrid cloud storage service that provides a local file system interface to S3. You can use it to store and retrieve files as objects in S3 using standard file protocols such as NFS and SMB. However, it does not support SFTP protocol, and it requires deploying a file gateway appliance on-premises or on EC2.
Option C is not cost-effective or scalable because it requires launching and managing an EC2 instance in a private subnet and setting up a VPN connection for the new partner. This would incur additional costs for the EC2 instance, the VPN connection, and the data transfer. It would also introduce complexity and security risks to the solution. Moreover, it would require running a cron job script on the EC2 instance to upload files to the S3 data lake, which is not efficient or reliable.
Option D is not cost-effective or scalable because it requires launching and managing multiple EC2 instances in a private subnet and placing a NLB in front of them. This would incur additional costs for the EC2 instances, the NLB, and the data transfer. It would also introduce complexity and security risks to the solution. Moreover, it would require running a cron job script on the EC2 instances to upload files to the S3 data lake, which is not efficient or reliable. References:
? What Is AWS Transfer Family?
? What Is Amazon S3 File Gateway?
? What Is Amazon EC2?
? [What Is Amazon Virtual Private Cloud?]
? [What Is a Network Load Balancer?]

NEW QUESTION 78
- (Topic 4)
A company has resources across multiple AWS Regions and accounts. A newly hired solutions architect discovers a previous employee did not provide details about the resources invent^. The solutions architect needs to build and map the relationship details of the various workloads across all accounts.
Which solution will meet these requirements in the MOST operationally efficient way?

A. Use AWS Systems Manager Inventory to generate a map view from the detailed view report.
B. Use AWS Step Functions to collect workload details Build architecture diagrams of theworkloads manually.
C. Use Workload Discovery on AWS to generate architecture diagrams of the workloads.
D. Use AWS X-Ray to view the workload details Build architecture diagrams with relationships

Answer: C

Explanation:
Workload Discovery on AWS (formerly called AWS Perspective) is a tool that visualizes AWS Cloud workloads. It maintains an inventory of the AWS resources across your accounts and Regions, mapping relationships between them, and displaying them in a web UI. It also allows you to query AWS Cost and Usage Reports, search for resources, save and export architecture diagrams, and more1. By using Workload Discovery on AWS, the solution can build and map the relationship details of the various workloads across all accounts with the least operational effort.
* A. Use AWS Systems Manager Inventory to generate a map view from the detailed view report. This solution will not meet the requirement of building and mapping the relationship details of the various workloads across all accounts, as AWS Systems Manager Inventory is a feature that collects metadata from your managed instances and stores it in a central Amazon S3 bucket. It does not provide a map view or architecture diagrams of the workloads2.
* B. Use AWS Step Functions to collect workload details Build architecture diagrams of the work-loads manually. This solution will not meet the requirement of the least operational effort, as it involves creating and managing state machines to orchestrate the workload details collection, and building architecture diagrams manually.
* D. Use AWS X-Ray to view the workload details Build architecture diagrams with relationships. This solution will not meet the requirement of the least operational effort, as it involves instrumenting your applications with X-Ray SDKs to collect workload details, and building architecture diagrams manually.
Reference URL: https://aws.amazon.com/solutions/implementations/workload-discovery- on-aws/

NEW QUESTION 80
- (Topic 4)
A company has created a multi-tier application for its ecommerce website. The website uses an Application Load Balancer that resides in the public subnets, a web tier in the public subnets, and a MySQL cluster hosted on Amazon EC2 instances in the private subnets. The MySQL database needs to retrieve product catalog and pricing information that is hosted on the internet by a third-party provider. A solutions architect must devise a strategy that maximizes security without increasing operational overhead. What should the solutions architect do to meet these requirements?

A. Deploy a NAT instance in the VP
B. Route all the internet-based traffic through the NAT instance.
C. Deploy a NAT gateway in the public subnet

D. Modify the private subnet route table to direct all internet-bound traffic to the NAT gateway.
E. Configure an internet gateway and attach it to the VP
F. Modify the private subnet route table to direct internet-bound traffic to the internet gateway.
G. Configure a virtual private gateway and attach it to the VP
H. Modify the private subnet route table to direct internet-bound traffic to the virtual private gateway.

Answer: B

Explanation:
To allow the MySQL database in the private subnets to access the internet without exposing it to the public, a NAT gateway is a suitable solution. A NAT gateway 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 instances. A NAT gateway resides in the public subnets and can handle high throughput of traffic with low latency. A NAT gateway is also a managed service that does not require any operational overhead. References:
? NAT Gateways
? NAT Gateway Pricing

NEW QUESTION 82
- (Topic 4)
A company previously migrated its data warehouse solution to AWS. The company also has an AWS Direct Connect connection. Corporate office users query the data warehouse using a visualization tool. The average size of a query returned by the data warehouse is 50 MB and each webpage sent by the visualization tool is approximately 500 KB. Result sets returned by the data warehouse are not cached.
Which solution provides the LOWEST data transfer egress cost for the company?

A. Host the visualization tool on premises and query the data warehouse directly over the internet.
B. Host the visualization tool in the same AWS Region as the data warehous
C. Access it over the internet.
D. Host the visualization tool on premises and query the data warehouse directly over a Direct Connect connection at a location in the same AWS Region.
E. Host the visualization tool in the same AWS Region as the data warehouse and access it over a Direct Connect connection at a location in the same Region.

Answer: D

Explanation:
https://aws.amazon.com/directconnect/pricing/ https://aws.amazon.com/blogs/aws/aws-data-transfer-prices-reduced/

NEW QUESTION 83
- (Topic 4)
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 needs to receive monthly email messages if any financial information is present in the employee data.
Which combination of steps should a solutions architect take to meet these requirements? (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 hierarchie
D. Export the data to Amazon S3 every month.
E. Configure Amazon fvlacie for the AWS accoun
F. Integrate Macie with Amazon EventBridge to send monthly events to AWS Lambda.
G. 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.
H. Configure Amazon Macie for the AWS account Integrate Macie with Amazon EventBridge to send monthly notifications through an Amazon Simple Notification Service (Amazon SNS) subscription.

Answer: BE

Explanation:
Generally, for building a hierarchical relationship model, a graph database such as Amazon Neptune is a better choice. In some cases, however, DynamoDB is a better choice for hierarchical data modeling because of its flexibility, security, performance, and scale. https://docs.aws.amazon.com/prescriptive-guidance/latest/dynamodb- hierarchical-data-model/introduction.html

NEW QUESTION 84
- (Topic 4)
A company runs a web application on Amazon EC2 instances in an Auto Scaling group that has a target group. The company desgned the application to work with session affinity (sticky sessions) for a better user experience.
The application must be available publicly over the internet as an endpoint_ A WAF must be applied to the endpoint for additional security. Session affinity (sticky sessions) must be configured on the endpoint
Which combination of steps will meet these requirements? (Select TWO)

A. Create a public Network Load Balancer Specify the application target group.
B. Create a Gateway Load Balancer Specify the application target group.
C. Create a public Application Load Balancer Specify the application target group.
D. Create a second target grou
E. Add Elastic IP addresses to the EC2 instances
F. Create a web ACL in AWS WAF Associate the web ACL with the endpoint

Answer: CE

Explanation:
C and E are the correct answers because they allow the company to create a public endpoint for its web application that supports session affinity (sticky sessions) and has a WAF applied for additional security. By creating a public Application Load Balancer, the company can distribute incoming traffic across multiple EC2 instances in an Auto Scaling group and specify the application target group. By creating a web ACL in AWS WAF and associating it with the Application Load

Balancer, the company can protect its web application from common web exploits. By enabling session stickiness on the Application Load Balancer, the company can ensure that subsequent requests from a user during a session are routed to the same target. References:
? Application Load Balancers
? AWS WAF
? Target Groups for Your Application Load Balancers
? How Application Load Balancer Works with Sticky Sessions

NEW QUESTION 86
- (Topic 4)
An ecommerce company stores terabytes of customer data in the AWS Cloud. The data contains personally identifiable information (Pll). The company wants to use the data in three applications. Only one of the applications needs to process the Pll. The Pll must be removed before the other two applications process the data.
Which solution will meet these requirements with the LEAST operational overhead?

A. Store the data in an Amazon DynamoDB tabl
B. Create a proxy application layer to intercept and process the data that each application requests.
C. Store the data in an Amazon S3 bucke
D. Process and transform the data by using S3 Object Lambda before returning the data to the requesting application.
E. Process the data and store the transformed data in three separate Amazon S3 buckets so that each application has its own custom datase
F. Point each application to its respectiveS3 bucket.
G. Process the data and store the transformed data in three separate Amazon DynamoDB tables so that each application has its own custom datase
H. Point each application to its respective DynamoDB table.

Answer: B

Explanation:
https://aws.amazon.com/blogs/aws/introducing-amazon-s3-object-lambda- use-your-code-to-process-data-as-it-is-being-retrieved-from-s3/
S3 Object Lambda is a new feature of Amazon S3 that enables customers to add their own code to process data retrieved from S3 before returning it to the application. By using S3 Object Lambda, the data can be processed and transformed in real-time, without the need to store multiple copies of the data in separate S3 buckets or DynamoDB tables.
In this case, the Pll can be removed from the data by the code added to S3 Object Lambda before returning the data to the two applications that do not need to process Pll. The one application that requires Pll can be pointed to the original S3 bucket where the Pll is still stored.
Using S3 Object Lambda is the simplest and most cost-effective solution, as it eliminates the need to maintain multiple copies of the same data in different buckets or tables, which can result in additional storage costs and operational overhead.

NEW QUESTION 90
- (Topic 4)
A company has data collection sensors at different locations. The data collection sensors stream a high volume of data to the company. The company wants to design a platform on AWS to ingest and process high-volume streaming data. The solution must be scalable and support data collection in near real time. The company must store the data in Amazon S3 for future reporting.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use Amazon Kinesis Data Firehose to deliver streaming data to Amazon S3.
B. Use AWS Glue to deliver streaming data to Amazon S3.
C. Use AWS Lambda to deliver streaming data and store the data to Amazon S3.
D. Use AWS Database Migration Service (AWS DMS) to deliver streaming data to Amazon S3.

Answer: A

Explanation:
To ingest and process high-volume streaming data with the least operational overhead, Amazon Kinesis Data Firehose is a suitable solution. Amazon Kinesis Data Firehose can capture, transform, and deliver streaming data to Amazon S3 or other destinations. Amazon Kinesis Data Firehose can scale automatically to match the throughput of the data and handle any amount of data. Amazon Kinesis Data Firehose is also a fully managed service that does not require any servers to provision or manage. References:
? What Is Amazon Kinesis Data Firehose?
? Amazon Kinesis Data Firehose Pricing

NEW QUESTION 92
- (Topic 4)
A global marketing company has applications that run in the ap-southeast-2 Region and the eu-west-1 Region. Applications that run in a VPC in eu-west-1 need to communicate securely with databases that run in a VPC in ap-southeast-2.
Which network design will meet these requirements?

A. Create a VPC peering connection between the eu-west-1 VPC and the ap-southeast-2 VP
B. Create an inbound rule in the eu-west-1 application security group that allows traffic from the database server IP addresses in the ap-southeast-2 security group.
C. Configure a VPC peering connection between the ap-southeast-2 VPC and the eu-west- 1 VP
D. Update the subnet route table
E. Create an inbound rule in the ap-southeast-2 database security group that references the security group ID of the application servers in eu-west-1.
F. Configure a VPC peering connection between the ap-southeast-2 VPC and the eu-west- 1 VP
G. Update the subnet route tables Create an inbound rule in the ap-southeast-2 database security group that allows traffic from the eu-west-1 application server IP addresses.
H. Create a transit gateway with a peering attachment between the eu-west-1 VPC and the ap-southeast-2 VP
I. After the transit gateways are properly peered and routing is configured, create an inbound rule in the database security group that references the security group ID of the application servers in eu-west-1.

Answer: C

Explanation:
"You cannot reference the security group of a peer VPC that's in a different Region. Instead, use the CIDR block of the peer VPC."

https://docs.aws.amazon.com/vpc/latest/peering/vpc-peering-security-groups.html

NEW QUESTION 94
- (Topic 4)
A company is running a photo hosting service in the us-east-1 Region. The service enables users across multiple countries to upload and view photos. Some photos are heavily viewed for months, and others are viewed for less than a week. The application allows uploads of up to 20 MB for each photo. The service uses the photo metadata to determine which photos to display to each user.
Which solution provides the appropriate user access MOST cost-effectively?

A. Store the photos in Amazon DynamoD
B. Turn on DynamoDB Accelerator (DAX) to cache frequently viewed items.
C. Store the photos in the Amazon S3 Intelligent-Tiering storage clas
D. Store the photo metadata and its S3 location in DynamoDB.
E. Store the photos in the Amazon S3 Standard storage clas
F. Set up an S3 Lifecycle policy to move photos older than 30 days to the S3 Standard-Infrequent Access (S3 Standard-IA) storage clas
G. Use the object tags to keep track of metadata.
H. Store the photos in the Amazon S3 Glacier storage clas
I. Set up an S3 Lifecycle policy to move photos older than 30 days to the S3 Glacier Deep Archive storage clas
J. Store the photo metadata and its S3 location in Amazon OpenSearch Service.

Answer: B

Explanation:
This solution provides the appropriate user access most cost-effectively
because it uses the Amazon S3 Intelligent-Tiering storage class, which automatically optimizes storage costs by moving data to the most cost-effective access tier when access patterns change, without performance impact or operational overhead1. This storage class is ideal for data with unknown, changing, or unpredictable access patterns, such as photos that are heavily viewed for months or less than a week. By storing the photo metadata and its S3 location in DynamoDB, the application can quickly query and retrieve the relevant photos for each user. DynamoDB is a fast, scalable, and fully managed NoSQL database service that supports key-value and document data models2.
References: 1: Amazon S3 Intelligent-Tiering Storage Class | AWS3, Overview section2: Amazon DynamoDB - NoSQL Cloud Database Service4, Overview section.

NEW QUESTION 98
- (Topic 4)
A company stores critical data in Amazon DynamoDB tables in the company's AWS account. An IT administrator accidentally deleted a DynamoDB table. The deletion caused a significant loss of data and disrupted the company's operations. The company wants to prevent this type of disruption in the future.
Which solution will meet this requirement with the LEAST operational overhead?

A. Configure a trail in AWS CloudTrai
B. Create an Amazon EventBridge rule for delete action
C. Create an AWS Lambda function to automatically restore deleted DynamoDBtables.
D. Create a backup and restore plan for the DynamoDB table
E. Recover the DynamoDB tables manually.
F. Configure deletion protection on the DynamoDB tables.
G. Enable point-in-time recovery on the DynamoDB tables.

Answer: C

Explanation:
Deletion protection is a feature of DynamoDB that prevents accidental deletion of tables. When deletion protection is enabled, you cannot delete a table unless you explicitly disable it first. This adds an extra layer of security and reduces the risk of data loss and operational disruption. Deletion protection is easy to enable and disable using the AWS Management Console, the AWS CLI, or the DynamoDB API. This solution has the least operational overhead, as you do not need to create, manage, or invoke any additional resources or services. References:
? Using deletion protection to protect your table
? Preventing Accidental Table Deletion in DynamoDB
? Amazon DynamoDB now supports table deletion protection

NEW QUESTION 103
- (Topic 4)
A company is designing a new web application that will run on Amazon EC2 Instances. The application will use Amazon DynamoDB for backend data storage. The application traffic will be unpredictable. T company expects that the application read and write throughput to the database will be moderate to high. The company needs to scale in response to application traffic.
Which DynamoDB table configuration will meet these requirements MOST cost-effectively?

A. Configure DynamoDB with provisioned read and write by using the DynamoDB Standard table clas
B. Set DynamoDB auto scaling to a maximum defined capacity.
C. Configure DynamoDB in on-demand mode by using the DynamoDB Standard table class.
D. Configure DynamoDB with provisioned read and write by using the DynamoDB Standard Infrequent Access (DynamoDB Standard-IA) table clas
E. Set DynamoDB auto scaling to a maximum defined capacity.
F. Configure DynamoDB in on-demand mode by using the DynamoDB Standard Infrequent Access (DynamoDB Standard-IA) table class.

Answer: B

Explanation:
The most cost-effective DynamoDB table configuration for the web application is to configure DynamoDB in on-demand mode by using the DynamoDB Standard table class. This configuration will allow the company to scale in response to application traffic and pay only for the read and write requests that the application performs on the table.
On-demand mode is a flexible billing option that can handle thousands of requests per second without capacity planning. On-demand mode automatically adjusts the table’s capacity based on the incoming traffic, and charges only for the read and write requests that are actually performed. On-demand mode is suitable for applications with unpredictable or variable workloads, or applications that prefer the ease of paying for only what they use1.

The DynamoDB Standard table class is the default and recommended table class for most workloads. The DynamoDB Standard table class offers lower throughput costs than the DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA) table class, and is more cost-effective for tables where throughput is the dominant cost. The DynamoDB Standard table class also offers the same performance, durability, and availability as the DynamoDB Standard-IA table class2. The other options are not correct because they are either not cost-effective or not suitable for the use case. Configuring DynamoDB with provisioned read and write by using the DynamoDB Standard table class, and setting DynamoDB auto scaling to a maximum defined capacity is not correct because this configuration requires manual estimation and management of the table's capacity, which adds complexity and cost to the solution. Provisioned mode is a billing option that requires users to specify the amount of read and write capacity units for their tables, and charges for the reserved capacity regardless of usage. Provisioned mode is suitable for applications with predictable or stable workloads, or applications that require finer-grained control over their capacity settings1. Configuring DynamoDB with provisioned read and write by using the DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA) table class, and setting DynamoDB auto scaling to a maximum defined capacity is not correct because this configuration is not cost-effective for tables with moderate to high throughput. The DynamoDB Standard-IA table class offers lower storage costs than the DynamoDB Standard table class, but higher throughput costs. The DynamoDB Standard-IA table class is optimized for tables where storage is the dominant cost, such as tables that store infrequently accessed data2. Configuring DynamoDB in on-demand mode by using the DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA) table class is not correct because this configuration is not cost- effective for tables with moderate to high throughput. As mentioned above, the DynamoDB Standard-IA table class has higher throughput costs than the DynamoDB Standard table class, which can offset the savings from lower storage costs.
References:
? Table classes - Amazon DynamoDB
? Read/write capacity mode - Amazon DynamoDB

NEW QUESTION 104
- (Topic 4)
A company runs multiple Amazon EC2 Linux instances in a VPC across two Availability Zones. The instances host applications that use a hierarchical directory structure. The applications need to read and write rapidly and concurrently to shared storage.
What should a solutions architect do to meet these requirements?

A. Create an Amazon S3 bucke
B. Allow access from all the EC2 instances in the VPC.
C. Create an Amazon Elastic File System (Amazon EFS) file syste
D. Mount the EFS file system from each EC2 instance.
E. Create a file system on a Provisioned IOPS SSD (102) Amazon Elastic Block Store (Amazon EBS) volum
F. Attach the EBS volume to all the EC2 instances.
G. Create file systems on Amazon Elastic Block Store (Amazon EBS) volumes that are attached to each EC2 instanc
H. Synchromze the EBS volumes across the different EC2 instances.

Answer: B

Explanation:
it allows the EC2 instances to read and write rapidly and concurrently to shared storage across two Availability Zones. Amazon EFS provides a scalable, elastic, and highly available file system that can be mounted from multiple EC2 instances. Amazon EFS supports high levels of throughput and IOPS, and consistent low latencies. Amazon EFS also supports NFSv4 lock upgrading and downgrading, which enables high levels of concurrency. References:
? Amazon EFS Features
? Using Amazon EFS with Amazon EC2

NEW QUESTION 105
- (Topic 4)
A company is deploying an application in three AWS Regions using an Application Load Balancer Amazon Route 53 will be used to distribute traffic between these Regions. Which Route 53 configuration should a solutions architect use to provide the MOST high- performing experience?

A. Create an A record with a latency policy.
B. Create an A record with a geolocation policy.
C. Create a CNAME record with a failover policy.
D. Create a CNAME record with a geoproximity policy.

Answer: A

Explanation:
To provide the most high-performing experience for the users of the application, a solutions architect should use a latency routing policy for the Route 53 A record. This policy allows Route 53 to route traffic to the AWS Region that provides the lowest possible latency for the users1. A latency routing policy can also improve the availability of the application, as Route 53 can automatically route traffic to another Region if the primary Region becomes unavailable2.
References:
? 1: https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing- policy.html#routing-policy-latency
? 2: https://aws.amazon.com/route53/faqs/#Latency_Based_Routing

NEW QUESTION 108
- (Topic 4)
A company has two VPCs that are located in the us-west-2 Region within the same AWS account. The company needs to allow network traffic between these VPCs. Approximately 500 GB of data transfer will occur between the VPCs each month.
What is the MOST cost-effective solution to connect these VPCs?

A. Implement AWS Transit Gateway to connect the VPC
B. Update the route tables of each VPC to use the transit gateway for inter-VPC communication.
C. Implement an AWS Site-to-Site VPN tunnel between the VPC
D. Update the route tables of each VPC to use the VPN tunnel for inter-VPC communication.
E. Set up a VPC peering connection between the VPC
F. Update the route tables of each VPC to use the VPC peering connection for inter-VPC communication.
G. Set up a 1 GB AWS Direct Connect connection between the VPC
H. Update the route tables of each VPC to use the Direct Connect connection for inter-VPC communication.

Answer: C

Explanation:
To connect two VPCs in the same Region within the same AWS account, VPC peering is the most cost-effective solution. VPC peering allows direct network traffic between the VPCs without requiring a gateway, VPN connection, or AWS Transit Gateway. VPC peering also does not incur any additional charges for data transfer between the VPCs.
References:
? What Is VPC Peering?
? VPC Peering Pricing

NEW QUESTION 112
- (Topic 4)
A company wants to rearchitect a large-scale web application to a serverless microservices architecture. The application uses Amazon EC2 instances and is written in Python.
The company selected one component of the web application to test as a microservice. The component supports hundreds of requests each second. The company wants to create and test the microservice on an AWS solution that supports Python. The solution must also scale automatically and require minimal infrastructure and minimal operational support.
Which solution will meet these requirements?

A. Use a Spot Fleet with auto scaling of EC2 instances that run the most recent Amazon Linux operating system.
B. Use an AWS Elastic Beanstalk web server environment that has high availability configured.
C. Use Amazon Elastic Kubernetes Service (Amazon EKS). Launch Auto Scaling groups of self-managed EC2 instances.
D. Use an AWS Lambda function that runs custom developed code.

Answer: D

Explanation:
AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. You can use Lambda to create and test microservices that are written in Python or other supported languages. Lambda scales automatically to handle the number of requests per second. You only pay for the compute time you consume. Lambda also integrates with other AWS services, such as Amazon API Gateway, Amazon S3, Amazon DynamoDB, and Amazon SQS, to enable event-driven architectures. Lambda has minimal infrastructure and operational overhead, as you do not need to manage servers, operating systems, patches, or scaling policies.
The other options are not serverless solutions and require more infrastructure and operational support. They also do not scale automatically to handle the number of requests per second. A Spot Fleet is a collection of EC2 instances that run on spare capacity at low prices. However, Spot Instances can be interrupted by AWS at any time, which can affect the availability and performance of your microservice. AWS Elastic Beanstalk is a service that automates the deployment and management of web applications on EC2 instances. However, you still need to provision, configure, and monitor the underlying EC2 instances and load balancers. Amazon EKS is a service that runs Kubernetes on AWS. However, you still need to create, configure, and manage the EC2 instances that form the Kubernetes cluster and nodes. You also need to install and update the Kubernetes software and tools. References:
? What is AWS Lambda?
? Building Lambda functions with Python
? Create a layer for a Lambda Python function
? AWS Lambda – Function in Python
? How do I call my AWS Lambda function from a local python script?

NEW QUESTION 116
- (Topic 4)
A company runs a real-time data ingestion solution on AWS. The solution consists of the most recent version of Amazon Managed Streaming for Apache Kafka (Amazon MSK). The solution is deployed in a VPC in private subnets across three Availability Zones.
A solutions architect needs to redesign the data ingestion solution to be publicly available over the internet. The data in transit must also be encrypted.
Which solution will meet these requirements with the MOST operational efficiency?

A. Configure public subnets in the existing VP
B. Deploy an MSK cluster in the public subnet
C. Update the MSK cluster security settings to enable mutual TLS authentication.
D. Create a new VPC that has public subnet
E. Deploy an MSK cluster in the public subnet
F. Update the MSK cluster security settings to enable mutual TLS authentication.
G. Deploy an Application Load Balancer (ALB) that uses private subnet
H. Configure an ALB security group inbound rule to allow inbound traffic from the VPC CIDR block for HTTPS protocol.
I. Deploy a Network Load Balancer (NLB) that uses private subnet
J. Configure an NLB listener for HTTPS communication over the internet.

Answer: A

Explanation:
The solution that meets the requirements with the most operational efficiency is to configure public subnets in the existing VPC and deploy an MSK cluster in the public subnets. This solution allows the data ingestion solution to be publicly available over the internet without creating a new VPC or deploying a load balancer. The solution also ensures that the data in transit is encrypted by enabling mutual TLS authentication, which requires both the client and the server to present certificates for verification. This solution leverages the public access feature of Amazon MSK, which is available for clusters running Apache Kafka 2.6.0 or later versions1.
The other solutions are not as efficient as the first one because they either create unnecessary resources or do not encrypt the data in transit. Creating a new VPC with public subnets would incur additional costs and complexity for managing network resources and routing. Deploying an ALB or an NLB would also add more costs and latency for the data ingestion solution. Moreover, an ALB or an NLB would not encrypt the data in transit by itself, unless they are configured with HTTPS listeners and certificates, which would require additional steps and maintenance. Therefore, these solutions are not optimal for the given requirements.
References:
? Public access - Amazon Managed Streaming for Apache Kafka

NEW QUESTION 117
- (Topic 4)
A company is running its production and nonproduction environment workloads in multiple AWS accounts. The accounts are in an organization in AWS Organizations. The company needs to design a solution that will prevent the modification of cost usage tags.
Which solution will meet these requirements?

A. Create a custom AWS Config rule to prevent tag modification except by authorized principals.
B. Create a custom trail in AWS CloudTrail to prevent tag modification
C. Create a service control policy (SCP) to prevent tag modification except by authonzed principals.
D. Create custom Amazon CloudWatch logs to prevent tag modification.

Answer: C

Explanation:
This solution meets the requirements because it uses SCPs to restrict the actions that can be performed on cost usage tags in the organization. SCPs are a type of organization policy that you can use to manage permissions in your organization. SCPs specify the maximum permissions for an organization, organizational unit (OU), or account. You can use SCPs to enforce consistent tag policies across your organization and prevent unauthorized or accidental changes to your tags. You can also create exceptions for authorized principals, such as administrators or auditors, who need to modify tags for legitimate purposes.
References:
? Service control policies (SCPs) - AWS Organizations
? Tag policies - AWS Organizations

NEW QUESTION 122
- (Topic 4)
A company has applications hosted on Amazon EC2 instances with IPv6 addresses. The applications must initiate communications with other external applications using the internet.
However, the company's security policy states that any external service cannot initiate a connection to the EC2 instances.
What should a solutions architect recommend to resolve this issue?

A. Create a NAT gateway and make it the destination of the subnet's route table.
B. Create an internet gateway and make it the destination of the subnet's route table
C. Create a virtual private gateway and make it the destination of the subnet's route table.
D. Create an egress-only internet gateway and make it the destination of the subnet's route table.

Answer: D

Explanation:
An egress-only internet gateway is a VPC component that allows outbound communication over IPv6 from instances in your VPC to the internet, and prevents the internet from initiating an IPv6 connection with your instances. This meets the company's security policy and requirements. To use an egress-only internet gateway, you need to add a route in the subnet's route table that routes IPv6 internet traffic (::/0) to the egress-only internet gateway.
Reference URLs:
1 https://docs.aws.amazon.com/vpc/latest/userguide/egress-only-internet-gateway.html
2 https://dev.to/aws-builders/what-is-an-egress-only-internet-gateways-in-aws-7gp
3 https://docs.aws.amazon.com/vpc/latest/userguide/route-table-options.html

NEW QUESTION 125
- (Topic 4)
A company runs a Java-based job on an Amazon EC2 instance. The job runs every hour and takes 10 seconds to run. The job runs on a scheduled interval and consumes 1 GB of memory. The CPU utilization of the instance is low except for short surges during which the job uses the maximum CPU available. The company wants to optimize the costs to run the job.
Which solution will meet these requirements?

A. Use AWS App2Container (A2C) to containerize the jo
B. Run the job as an Amazon Elastic Container Service (Amazon ECS) task on AWS Fargate with 0.5 virtual CPU (vCPU) and 1 GB of memory.
C. Copy the code into an AWS Lambda function that has 1 GB of memor
D. Create an Amazon EventBridge scheduled rule to run the code each hour.
E. Use AWS App2Container (A2C) to containerize the jo
F. Install the container in the existing Amazon Machine Image (AMI). Ensure that the schedule stops the container when the task finishes.
G. Configure the existing schedule to stop the EC2 instance at the completion of the job and restart the EC2 instance when the next job starts.

Answer: B

Explanation:
AWS Lambda is a serverless compute service that allows you to run code without provisioning or managing servers. You can create Lambda functions using various languages, including Java, and specify the amount of memory and CPU allocated to your function. Lambda charges you only for the compute time you consume, which is calculated based on the number of requests and the duration of your code execution. You can use Amazon EventBridge to trigger your Lambda function on a schedule, such as every hour, using cron or rate expressions. This solution will optimize the costs to run the job, as you will not pay for any idle time or unused resources, unlike running the job on an EC2 instance. References: 1: AWS Lambda - FAQs2, General Information section2: Tutorial: Schedule AWS Lambda functions using EventBridge3, Introduction section3: Schedule expressions using rate or cron - AWS Lambda4, Introduction section.

NEW QUESTION 126
- (Topic 4)
A company needs to provide customers with secure access to its data. The company processes customer data and stores the results in an Amazon S3 bucket.
All the data is subject to strong regulations and security requirements. The data must be encrypted at rest. Each customer must be able to access only their data from their AWS account. Company employees must not be able to access the data.
Which solution will meet these requirements?

A. Provision an AWS Certificate Manager (ACM) certificate for each custome
B. Encrypt the data client-sid
C. In the private certificate policy, deny access to the certificate for all principals except an 1AM role that the customer provides.
D. Provision a separate AWS Key Management Service (AWS KMS) key for each custome
E. Encrypt the data server-sid
F. In the S3 bucket policy, deny decryption of data for all principals except an 1AM role that the customer provides.
G. Provision a separate AWS Key Management Service (AWS KMS) key for each custome
H. Encrypt the data server-sid
I. In each KMS key policy, deny decryption of data for all principals except an 1AM role that the customer provides.

J. Provision an AWS Certificate Manager (ACM) certificate for each custome
K. Encrypt the data client-sid
L. In the public certificate policy, deny access to the certificate for all principals except an 1AM role that the customer provides.

Answer: C

Explanation:
The correct solution is to provision a separate AWS KMS key for each customer and encrypt the data server-side. This way, the company can use the S3 encryption feature to protect the data at rest and delegate the control of the encryption keys to the customers. The customers can then use their own IAM roles to access and decrypt their data. The company employees will not be able to access the data because they are not authorized by the KMS key policies. The other options are incorrect because:
? Option A and D are using ACM certificates to encrypt the data client-side. This is
not a recommended practice for S3 encryption because it adds complexity and overhead to the encryption process. Moreover, the company will have to manage the certificates and their policies for each customer, which is not scalable and secure.
? Option B is using a separate KMS key for each customer, but it is using the S3
bucket policy to control the decryption access. This is not a secure solution because the bucket policy applies to the entire bucket, not to individual objects. Therefore, the customers will be able to access and decrypt each other's data if they have the permission to list the bucket contents. The bucket policy also overrides the KMS key policy, which means the company employees can access the data if they have the permission to use the KMS key.
References:
? S3 encryption
? KMS key policies
? ACM certificates

NEW QUESTION 129
- (Topic 4)
A social media company runs its application on Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is the origin for an Amazon CloudFront distribution. The application has more than a billion images stored in an Amazon S3 bucket and processes thousands of images each second. The company wants to resize the images dynamically and serve appropriate formats to clients.
Which solution will meet these requirements with the LEAST operational overhead?

A. Install an external image management library on an EC2 instanc
B. Use the imagemanagement library to process the images.
C. Create a CloudFront origin request polic
D. Use the policy to automatically resize images and to serve the appropriate format based on the User-Agent HTTP header in the request.
E. Use a Lambda@Edge function with an external image management librar
F. Associate the Lambda@Edge function with the CloudFront behaviors that serve the images.
G. Create a CloudFront response headers polic
H. Use the policy to automatically resize images and to serve the appropriate format based on the User-Agent HTTP header in the request.

Answer: C

Explanation:
To resize images dynamically and serve appropriate formats to clients, a Lambda@Edge function with an external image management library can be used. Lambda@Edge allows running custom code at the edge locations of CloudFront, which can process the images on the fly and optimize them for different devices and browsers. An external image management library can provide various image manipulation and optimization features. References:
? Lambda@Edge
? Resizing Images with Amazon CloudFront & Lambda@Edge

NEW QUESTION 133
- (Topic 4)
A company manages AWS accounts in AWS Organizations. AWS 1AM Identity Center (AWS Single Sign-On) and AWS Control Tower are configured for the accounts. The company wants to manage multiple user permissions across all the accounts.
The permissions will be used by multiple 1AM users and must be split between the developer and administrator teams. Each team requires different permissions. The company wants a solution that includes new users that are hired on both teams.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create individual users in 1AM Identity Center (or each accoun
B. Create separate developer and administrator groups in 1AM Identity Cente
C. Assign the users to the appropriate groups Create a custom 1AM policy for each group to set fine-grained permissions.
D. Create individual users in 1AM Identity Center for each accoun
E. Create separate developer and administrator groups in 1AM Identity Cente
F. Assign the users to the appropriate group
G. Attach AWS managed 1AM policies to each user as needed for fine- grained permissions.
H. Create individual users in 1AM Identity Center Create new developer and administrator groups in 1AM Identity Cente
I. Create new permission sets that include the appropriate 1AM policies for each grou
J. Assign the new groups to the appropriate accounts Assign the new permission sets to the new groups When new users are hired, add them to the appropriate group.
K. Create individual users in 1AM Identity Cente
L. Create new permission sets that include the appropriate 1AM policies for each use
M. Assign the users to the appropriate account
N. Grant additional 1AM permissions to the users from within specific account
O. When new users are hired, add them to 1AM Identity Center and assign them to the accounts.

Answer: C

Explanation:
This solution meets the requirements with the least operational overhead because it leverages the features of IAM Identity Center and AWS Control Tower to centrally manage multiple user permissions across all the accounts. By creating new groups and permission sets, the company can assign fine-grained permissions to the developer and administrator teams based on their roles and responsibilities. The permission sets are applied to the groups at the organization level, so they are automatically inherited by all the accounts in the organization. When new users are hired, the company only needs to add them to the appropriate group in IAM Identity Center, and they will automatically get the permissions assigned to that group. This simplifies the user management and reduces

the manual effort of assigning permissions to each user individually.
References:
? Managing access to AWS accounts and applications
? Managing permissions sets
? Managing groups

NEW QUESTION 136
- (Topic 4)
A company uses high concurrency AWS Lambda functions to process a constantly increasing number of messages in a message queue during marketing events. The Lambda functions use CPU intensive code to process the messages. The company wants to reduce the compute costs and to maintain service latency for its customers.
Which solution will meet these requirements?

A. Configure reserved concurrency for the Lambda function
B. Decrease the memory allocated to the Lambda functions.
C. Configure reserved concurrency for the Lambda function
D. Increase the memory according to AWS Compute Optimizer recommendations.
E. Configure provisioned concurrency for the Lambda function
F. Decrease the memory allocated to the Lambda functions.
G. Configure provisioned concurrency for the Lambda function
H. Increase the memory according to AWS Compute Optimizer recommendations.

Answer: D

Explanation:
The company wants to reduce the compute costs and maintain service latency for its Lambda functions that process a constantly increasing number of messages in a message queue. The Lambda functions use CPU intensive code to process the messages. To meet these requirements, a solutions architect should recommend the following solution:
? Configure provisioned concurrency for the Lambda functions. Provisioned
concurrency is the number of pre-initialized execution environments that are allocated to the Lambda functions. These execution environments are prepared to respond immediately to incoming function requests, reducing the cold start latency. Configuring provisioned concurrency also helps to avoid throttling errors due to reaching the concurrency limit of the Lambda service.
? Increase the memory according to AWS Compute Optimizer recommendations.
AWS Compute Optimizer is a service that provides recommendations for optimal AWS resource configurations based on your utilization data. By increasing the memory allocated to the Lambda functions, you can also increase the CPU power and improve the performance of your CPU intensive code. AWS Compute Optimizer can help you find the optimal memory size for your Lambda functions based on your workload characteristics and performance goals.
This solution will reduce the compute costs by avoiding unnecessary over-provisioning of memory and CPU resources, and maintain service latency by using provisioned concurrency and optimal memory size for the Lambda functions.
References:
? Provisioned Concurrency
? AWS Compute Optimizer

NEW QUESTION 138
- (Topic 4)
A company runs applications on AWS that connect to the company's Amazon RDS database. The applications scale on weekends and at peak times of the year. The company wants to scale the database more effectively for its applications that connect to the database.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use Amazon DynamoDB with connection pooling with a target group configuration for the databas
B. Change the applications to use the DynamoDB endpoint.
C. Use Amazon RDS Proxy with a target group for the databas
D. Change the applications to use the RDS Proxy endpoint.
E. Use a custom proxy that runs on Amazon EC2 as an intermediary to the databas
F. Change the applications to use the custom proxy endpoint.
G. Use an AWS Lambda function to provide connection pooling with a target group configuration for the databas
H. Change the applications to use the Lambda function.

Answer: B

Explanation:
Amazon RDS Proxy is a fully managed, highly available database proxy for Amazon Relational Database Service (RDS) that makes applications more scalable, more resilient to database failures, and more secure1. RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability2. RDS Proxy also reduces failover times for Aurora and RDS databases by up to 66% and enables IAM authentication and Secrets Manager integration for database access1. RDS Proxy can be enabled for most applications with no code changes2.

NEW QUESTION 139
- (Topic 4)
A serverless application uses Amazon API Gateway. AWS Lambda, and Amazon DynamoDB. The Lambda function needs permissions to read and write to the DynamoDB table.
Which solution will give the Lambda function access to the DynamoDB table MOST securely?

A. Create an 1AM user with programmatic access to the Lambda functio
B. Attach a policy to the user that allows read and write access to the DynamoDB tabl
C. Store the access_key_id and secret_access_key parameters as part of the Lambda environment variable
D. Ensure that other AWS users do not have read and write access to the Lambda function configuration
E. Create an 1AM role that includes Lambda as a trusted servic
F. Attach a policy to the role that allows read and write access to the DynamoDB tabl
G. Update the configuration of the Lambda function to use the new role as the execution role.
H. Create an 1AM user with programmatic access to the Lambda functio
I. Attach a policy to the user that allows read and write access to the DynamoDB tabl

J. Store the access_key_id and secret_access_key parameters in AWS Systems Manager Parameter Store as secure string parameter
K. Update the Lambda function code to retrieve the secure string parameters before connecting to the DynamoDB table.
L. Create an 1AM role that includes DynamoDB as a trusted servic
M. Attach a policy to the role that allows read and write access from the Lambda functio
N. Update the code of the Lambda function to attach to the new role as an execution role.

Answer: B

Explanation:
Option B suggests creating an IAM role that includes Lambda as a trusted service, meaning the role is specifically designed for Lambda functions. The role should have a policy attached to it that grants the required read and write access to the DynamoDB table.

NEW QUESTION 142
- (Topic 4)
A company wants to run its payment application on AWS The application receives payment notifications from mobile devices Payment notifications require a basic validation before they are sent for further processing
The backend processing application is long running and requires compute and memory to be adjusted The company does not want to manage the infrastructure
Which solution will meet these requirements with the LEAST operational overhead?

A. Create an Amazon Simple Queue Service (Amazon SQS) queue Integrate the queue with an Amazon EventBndge rule to receive payment notifications from mobile devices Configure the rule to validate payment notifications and send the notifications to the backend application Deploy the backend application on Amazon Elastic Kubernetes Service (Amazon EKS) Anywhere Create a standalone cluster
B. Create an Amazon API Gateway API Integrate the API with anAWS Step Functions state machine to receive payment notifications from mobile devices Invoke the state machine to validate payment notifications and send the notifications to the backend application Deploy the backend application on Amazon Elastic Kubernetes Sen/ice (Amazon EKS). Configure an EKS cluster with self-managed nodes.
C. Create an Amazon Simple Queue Sen/ice (Amazon SQS) queue Integrate the queue with an Amazon EventBridge rule to receive payment notifications from mobile devices Configure the rule to validate payment notifications and send the notifications to the backend application Deploy the backend application on Amazon EC2 Spot Instances Configure a Spot Fleet with a default allocation strategy.
D. Create an Amazon API Gateway API Integrate the API with AWS Lambda to receive payment notifications from mobile devices Invoke a Lambda function to validate payment notifications and send the notifications to the backend application Deploy the backend application on Amazon Elastic Container Service (Amazon ECS). Configure Amazon ECS with an AWS Fargate launch type.

Answer: D

Explanation:
This option is the best solution because it allows the company to run its payment application on AWS with minimal operational overhead and infrastructure management. By using Amazon API Gateway, the company can create a secure and scalable API to receive payment notifications from mobile devices. By using AWS Lambda, the company can run a serverless function to validate the payment notifications and send them to the backend application. Lambda handles the provisioning, scaling, and security of the function, reducing the operational complexity and cost. By using Amazon ECS with AWS Fargate, the company can run the backend application on a fully managed container service that scales the compute resources automatically and does not require any EC2 instances to manage. Fargate allocates the right amount of CPU and memory for each container and adjusts them as needed.
* A. Create an Amazon Simple Queue Service (Amazon SQS) queue Integrate the queue with an Amazon EventBndge rule to receive payment notifications from mobile devices Configure the rule to validate payment notifications and send the notifications to the backend application Deploy the backend application on Amazon Elastic Kubernetes Service (Amazon EKS) Anywhere Create a standalone cluster. This option is not optimal because it requires the company to manage the Kubernetes cluster that runs the backend application. Amazon EKS Anywhere is a deployment option that allows the company to create and operate Kubernetes clusters on-premises or in other environments outside AWS. The company would need to provision, configure, scale, patch, and monitor the cluster nodes, which can increase the operational overhead and complexity. Moreover, the company would need to ensure the connectivity and security between the AWS services and the EKS Anywhere cluster, which can also add challenges and risks.
* B. Create an Amazon API Gateway API Integrate the API with anAWS Step Functions state ma-chine to receive payment notifications from mobile devices Invoke the state machine to validate payment notifications and send the notifications to the backend application Deploy the backend application on Amazon Elastic Kubernetes Sen/ice (Amazon EKS). Configure an EKS cluster with self-managed nodes. This option is not ideal because it requires the company to manage the EC2 instances that host the Kubernetes cluster that runs the backend application. Amazon EKS is a fully managed service that runs Kubernetes on AWS, but it still requires the company to manage the worker nodes that run the containers. The company would need to provision, configure, scale, patch, and monitor the EC2 instances, which can increase the operational overhead and infrastructure costs. Moreover, using AWS Step Functions to validate the payment notifications may be unnecessary and complex, as the validation logic can be implemented in a simpler way with Lambda or other services.
* C. Create an Amazon Simple Queue Sen/ice (Amazon SQS) queue Integrate the queue with an Amazon EventBridge rule to receive payment notifications from mobile devices Configure the rule to validate payment notifications and send the notifications to the backend application Deploy the backend application on Amazon EC2 Spot Instances Configure a Spot Fleet with a default al-location strategy. This option is not cost-effective because it requires the company to manage the EC2 instances that run the backend application. The company would need to provision, configure, scale, patch, and monitor the EC2 instances, which can increase the operational overhead and infrastructure costs. Moreover, using Spot Instances can introduce the risk of interruptions, as Spot Instances are reclaimed by AWS when the demand for On-Demand Instances increases. The company would need to handle the interruptions gracefully and ensure the availability and reliability of the backend application.
References:
? 1 Amazon API Gateway - Amazon Web Services
? 2 AWS Lambda - Amazon Web Services
? 3 Amazon Elastic Container Service - Amazon Web Services
? 4 AWS Fargate - Amazon Web Services

NEW QUESTION 143
.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Solution-Architect-Associate Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-Solution-Architect-Associate Product From:

https://www.2passeasy.com/dumps/AWS-Solution-Architect-Associate/

Money Back Guarantee

AWS-Solution-Architect-Associate Practice Exam Features:

* AWS-Solution-Architect-Associate Questions and Answers Updated Frequently

* AWS-Solution-Architect-Associate Practice Questions Verified by Expert Senior Certified Staff

* AWS-Solution-Architect-Associate Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AWS-Solution-Architect-Associate Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year