Amazon

Exam Questions AWS-Solution-Architect-Associate

Amazon AWS Certified Solutions Architect - Associate

NEW QUESTION 1
- (Topic 4)
A solutions architect needs to design the architecture for an application that a vendor provides as a Docker container image The container needs 50 GB of storage available for temporary files The infrastructure must be serverless.
Which solution meets these requirements with the LEAST operational overhead?

A. Create an AWS Lambda function that uses the Docker container image with an Amazon S3 mounted volume that has more than 50 GB of space
B. Create an AWS Lambda function that uses the Docker container image with an Amazon Elastic Block Store (Amazon EBS) volume that has more than 50 GB of space
C. Create an Amazon Elastic Container Service (Amazon ECS) cluster that uses the AWS Fargate launch type Create a task definition for the container image with an Amazon Elastic File System (Amazon EFS) volum
D. Create a service with that task definition.
E. Create an Amazon Elastic Container Service (Amazon ECS) cluster that uses the Amazon EC2 launch type with an Amazon Elastic Block Store (Amazon EBS) volume that has more than 50 GB of space Create a task definition for the container imag
F. Create a service with that task definition.

Answer: C

Explanation:
The AWS Fargate launch type is a serverless way to run containers on Amazon ECS,
without having to manage any underlying infrastructure. You only pay for the resources required to run your containers, and AWS handles the provisioning, scaling, and security of the cluster. Amazon EFS is a fully managed, elastic, and scalable file system that can be mounted to multiple containers, and provides high availability and durability. By using AWS Fargate and Amazon EFS, you can run your Docker container image with 50 GB of storage available for temporary files, with the least operational overhead. This solution meets the requirements of the question.
References:
? AWS Fargate
? Amazon Elastic File System
? Using Amazon EFS file systems with Amazon ECS

NEW QUESTION 2
- (Topic 4)
A company has an on-premises MySQL database that handles transactional data. The company is migrating the database to the AWS Cloud. The migrated database must maintain compatibility with the company's applications that use the database. The migrated database also must scale automatically during periods of increased demand.
Which migration solution will meet these requirements?

A. Use native MySQL tools to migrate the database to Amazon RDS for MySQ
B. Configureelastic storage scaling.
C. Migrate the database to Amazon Redshift by using the mysqldump utilit
D. Turn on Auto Scaling for the Amazon Redshift cluster.
E. Use AWS Database Migration Service (AWS DMS) to migrate the database to Amazon Auror
F. Turn on Aurora Auto Scaling.
G. Use AWS Database Migration Service (AWS DMS) to migrate the database to Amazon DynamoD
H. Configure an Auto Scaling policy.

Answer: C

Explanation:
To migrate a MySQL database to AWS with compatibility and scalability, Amazon Aurora is a suitable option. Aurora is compatible with MySQL and can scale automatically with Aurora Auto Scaling. AWS Database Migration Service (AWS DMS) can be used to migrate the database from on-premises to Aurora with minimal downtime. References:
? What Is Amazon Aurora?
? Using Amazon Aurora Auto Scaling with Aurora Replicas
? What Is AWS Database Migration Service?

NEW QUESTION 3
- (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 4
- (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 5
- (Topic 4)
A company is moving its data and applications to AWS during a multiyear migration project. The company wants to securely access data on Amazon S3 from the company's AWS Region and from the company's on-premises location. The data must not traverse the internet. The company has established an AWS Direct Connect connection between its Region and its on-premises location
Which solution will meet these requirements?

A. Create gateway endpoints for Amazon S3. Use the gateway endpoints to securely access the data from the Region and the on-premises location.
B. Create a gateway in AWS Transit Gateway to access Amazon S3 securely from the Region and the on-premises location.
C. Create interface endpoints for Amazon S3_ Use the interface endpoints to securely access the data from the Region and the on-premises location.
D. Use an AWS Key Management Service (AWS KMS) key to access the data securely from the Region and the on-premises location.

Answer: B

Explanation:
A gateway endpoint is a gateway that is a target for a specified route in your route table, used for traffic destined to a supported AWS service1. Amazon S3 does not support gateway endpoints, only interface endpoints2. Therefore, option A is incorrect.
An interface endpoint is an elastic network interface with a private IP address that serves as an entry point for traffic destined to a supported service1. An interface endpoint can provide secure access to Amazon S3 from within the Region, but not from the on-premises location. Therefore, option C is incorrect.
AWS Key Management Service (AWS KMS) is a service that allows you to create and manage encryption keys to protect your data3. AWS KMS does not provide a way to access data on Amazon S3 without traversing the internet. Therefore, option D is incorrect. AWS Transit Gateway is a service that enables you to connect your Amazon Virtual Private Clouds (VPCs) and your on-premises networks to a single gateway. You can create a gateway in AWS Transit Gateway to access Amazon S3 securely from both the Region and the on-premises location using AWS Direct Connect. Therefore, option B is correct.

NEW QUESTION 6
- (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 CloudFront 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 CloudFront 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 7
- (Topic 4)
A company has an application that uses Docker containers in its local data center The application runs on a container host that stores persistent data in a volume on the host. The container instances use the stored persistent data.
The company wants to move the application to a fully managed service because the company does not want to manage any servers or storage infrastructure.
Which solution will meet these requirements?

A. Use Amazon Elastic Kubernetes Service (Amazon EKS) with self-managed node
B. Create an Amazon Elastic Block Store (Amazon EBS) volume attached to an Amazon EC2 instanc
C. Use the EBS volume as a persistent volume mounted in the containers.
D. Use Amazon Elastic Container Service (Amazon ECS) with an AWS Fargate launch typ
E. Create an Amazon Elastic File System (Amazon EFS) volum
F. Add the EFS volumeas a persistent storage volume mounted in the containers.
G. Use Amazon Elastic Container Service (Amazon ECS) with an AWS Fargate launch typ
H. Create an Amazon S3 bucke
I. Map the S3 bucket as a persistent storage volume mounted in the containers.
J. Use Amazon Elastic Container Service (Amazon ECS) with an Amazon EC2 launch typ
K. Create an Amazon Elastic File System (Amazon EFS) volum
L. Add the EFS volume as a persistent storage volume mounted in the containers.

Answer: B

Explanation:
This solution meets the requirements because it allows the company to move the application to a fully managed service without managing any servers or storage infrastructure. AWS Fargate is a serverless compute engine for containers that runs the Amazon ECS tasks. With Fargate, the company does not need to provision, configure, or scale clusters of virtual machines to run containers. Amazon EFS is a fully managed file system that can be accessed by multiple containers concurrently. With EFS, the company does not need to provision and manage storage capacity. EFS provides a simple interface to create and configure file systems quickly and easily. The company can use the EFS volume as a persistent storage volume mounted in the containers to store the persistent data. The company can also use the EFS mount helper to simplify the mounting process. References: Amazon ECS on AWS Fargate, Using Amazon EFS file systems with Amazon ECS, Amazon EFS mount helper.

NEW QUESTION 8
- (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 9
- (Topic 4)
A company needs to configure a real-time data ingestion architecture for its application. The company needs an API. a process that transforms data as the data is streamed, and a storage solution for the data.
Which solution will meet these requirements with the LEAST operational overhead?

A. Deploy an Amazon EC2 instance to host an API that sends data to an Amazon Kinesis data strea
B. Create an Amazon Kinesis Data Firehose delivery stream that uses the Kinesis data stream as a data sourc
C. Use AWS Lambda functions to transform the dat
D. Use the Kinesis Data Firehose delivery stream to send the data to Amazon S3.
E. Deploy an Amazon EC2 instance to host an API that sends data to AWS Glu
F. Stop source/destination checking on the EC2 instanc
G. Use AWS Glue to transform the data and to send the data to Amazon S3.
H. Configure an Amazon API Gateway API to send data to an Amazon Kinesis datastrea
I. Create an Amazon Kinesis Data Firehose delivery stream that uses the Kinesis data stream as a data sourc

J. Use AWS Lambda functions to transform the dat
K. Use the Kinesis Data Firehose delivery stream to send the data to Amazon S3.
L. Configure an Amazon API Gateway API to send data to AWS Glu
M. Use AWS Lambda functions to transform the dat
N. Use AWS Glue to send the data to Amazon S3.

Answer: C

Explanation:
It uses Amazon Kinesis Data Firehose which is a fully managed service for delivering real-time streaming data to destinations such as Amazon S3. This service requires less operational overhead as compared to option A, B, and D. Additionally, it also uses Amazon API Gateway which is a fully managed service for creating, deploying, and managing APIs. These services help in reducing the operational overhead and automating the data ingestion process.

NEW QUESTION 10
- (Topic 4)
A company has a multi-tier payment processing application that is based on virtual machines (VMs). The communication between the tiers occurs asynchronously through a third-party middleware solution that guarantees exactly-once delivery.
The company needs a solution that requires the least amount of infrastructure management. The solution must guarantee exactly-once delivery for application messaging
Which combination of actions will meet these requirements? (Select TWO.)

A. Use AWS Lambda for the compute layers in the architecture.
B. Use Amazon EC2 instances for the compute layers in the architecture.
C. Use Amazon Simple Notification Service (Amazon SNS) as the messaging component between the compute layers.
D. Use Amazon Simple Queue Service (Amazon SQS) FIFO queues as the messaging component between the compute layers.
E. Use containers that are based on Amazon Elastic Kubemetes Service (Amazon EKS) for the compute layers in the architecture.

Answer: AD

Explanation:
This solution meets the requirements because it requires the least amount of infrastructure management and guarantees exactly-once delivery for application messaging. AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. You only pay for the compute time you consume. Lambda scales automatically with the size of your workload. Amazon SQS FIFO queues are designed to ensure that messages are processed exactly once, in the exact order that they are sent. FIFO queues have high availability and deliver messages in a strict first-in, first-out order. You can use Amazon SQS to decouple and scale microservices, distributed systems, and serverless applications. References: AWS Lambda, Amazon SQS FIFO queues

NEW QUESTION 10
- (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 13
- (Topic 4)
A solutions architect is designing a highly available Amazon ElastiCache for Redis based solution. The solutions architect needs to ensure that failures do not result in performance degradation or loss of data locally and within an AWS Region. The solution needs to provide high availability at the node level and at the Region level.
Which solution will meet these requirements?

A. Use Multi-AZ Redis replication groups with shards that contain multiple nodes.
B. Use Redis shards that contain multiple nodes with Redis append only files (AOF) tured on.
C. Use a Multi-AZ Redis cluster with more than one read replica in the replication group.
D. Use Redis shards that contain multiple nodes with Auto Scaling turned on.

Answer: A

Explanation:
This answer is correct because it provides high availability at the node level and at the Region level for the ElastiCache for Redis solution. A Multi-AZ Redis replication group consists of a primary cluster and up to five read replica clusters, each in a different Availability Zone. If the primary cluster fails, one of the read replicas is automatically promoted to be the new primary cluster. A Redis replication group with shards enables partitioning of the data across multiple nodes, which increases the scalability and performance of the solution. Each shard can have one or more replicas to provide redundancy and read scaling.
References:
? https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/AutoFailover.html
? https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/Shards.html

NEW QUESTION 15
- (Topic 4)
A financial company needs to handle highly sensitive data The company will store the data in an Amazon S3 bucket The company needs to ensure that the data is encrypted in transit and at rest The company must manage the encryption keys outside the AWS Cloud
Which solution will meet these requirements?

A. Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) customer managed key
B. Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) AWS managed key
C. Encrypt the data in the S3 bucket with the default server-side encryption (SSE)
D. Encrypt the data at the company's data center before storing the data in the S3 bucket

Answer: D

Explanation:
This option is the only solution that meets the requirements because it allows the company to encrypt the data with its own encryption keys and tools outside the AWS Cloud. By encrypting the data at the company’s data center before storing the data in the S3 bucket, the company can ensure that the data is encrypted in transit and at rest, and that the company has full control over the encryption keys and processes. This option also avoids the need to use any AWS encryption services or features, which may not be compatible with the company’s security policies or compliance standards.
* A. Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) customer managed key. This option does not meet the requirements because it does not allow the company to manage the encryption keys outside the AWS Cloud. Although the company can create and use its own customer managed key in AWS KMS, the key is still stored and managed by AWS KMS, which is a service within the AWS Cloud. Moreover, the company still needs to use the AWS encryption features and APIs to encrypt and decrypt the data in the S3 bucket, which may not be compatible with the company’s security policies or compliance standards.
* B. Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) AWS managed key. This option does not meet the requirements because it does not allow the company to manage the encryption keys outside the AWS Cloud. In this option, the company uses the default AWS managed key in AWS KMS, which is created and managed by AWS on behalf of the company. The company has no control over the key rotation, deletion, or recovery policies. Moreover, the company still needs to use the AWS encryption features and APIs to encrypt and decrypt the data in the S3 bucket, which may not be compatible with the company’s security policies or compliance standards.
* C. Encrypt the data in the S3 bucket with the default server-side encryption (SSE). This option does not meet the requirements because it does not allow the company to manage the encryption keys outside the AWS Cloud. In this option, the company uses the default server-side encryption with Amazon S3 managed keys (SSE-S3), which is applied to every bucket in Amazon S3. The company has no visibility or control over the encryption keys, which are managed by Amazon S3. Moreover, the company still needs to use the AWS encryption features and APIs to encrypt and decrypt the data in the S3 bucket, which may not be compatible with the company’s security policies or compliance standards. References:
? 1 Protecting data with encryption - Amazon Simple Storage Service
? 2 Protecting data with server-side encryption - Amazon Simple Storage Service
? 3 Protecting data by using client-side encryption - Amazon Simple Storage Service
? 4 AWS Key Management Service Concepts - AWS Key Management Service

NEW QUESTION 17
- (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 20
- (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 25
- (Topic 4)
A company has an on-premises server that uses an Oracle database to process and store customer information The company wants to use an AWS database service to achieve higher availability and to improve application performance. The company also wants to offload reporting from its primary database system.
Which solution will meet these requirements in the MOST operationally efficient way?

A. Use AWS Database Migration Service (AWS DMS) to create an Amazon RDS DB instance in multiple AWS Regions Point the reporting functions toward a separate DB instance from the primary DB instance.
B. Use Amazon RDS in a Single-AZ deployment to create an Oracle database Create a read replica in the same zone as the primary DB instanc
C. Direct the reporting functions to the read replica.
D. Use Amazon RDS deployed in a Multi-AZ cluster deployment to create an Oracle database Direct the reporting functions to use the reader instance in the cluster deployment
E. Use Amazon RDS deployed in a Multi-AZ instance deployment to create an Amazon Aurora databas
F. Direct the reporting functions to the reader instances.

Answer: D

Explanation:
Amazon Aurora is a fully managed relational database that is compatible with MySQL and PostgreSQL. It provides up to five times better performance than MySQL and
up to three times better performance than PostgreSQL. It also provides high availability and durability by replicating data across multiple Availability Zones and continuously backing up data to Amazon S31. By using Amazon RDS deployed in a Multi-AZ instance deployment
to create an Amazon Aurora database, the solution can achieve higher availability and improve application performance.
Amazon Aurora supports read replicas, which are separate instances that share the same underlying storage as the primary instance. Read replicas can be used to offload read-only queries from the primary instance and improve performance. Read replicas can also be used for reporting functions2. By directing the reporting functions to the reader instances, the solution can offload reporting from its primary database system.
* A. Use AWS Database Migration Service (AWS DMS) to create an Amazon RDS DB instance in multiple AWS Regions Point the reporting functions toward a separate DB instance from the pri-mary DB instance. This solution will not meet the requirement of using an AWS database service, as AWS DMS is a service that helps users migrate databases to AWS, not a database service itself. It also involves creating multiple DB instances in different Regions, which may increase complexity and cost.
* B. Use Amazon RDS in a Single-AZ deployment to create an Oracle database Create a read replica in the same zone as the primary DB instance. Direct the reporting functions to the read replica. This solution will not meet the requirement of achieving higher availability, as a Single-AZ deployment does not provide failover protection in case of an Availability Zone outage. It also involves using Oracle as the database engine, which may not provide better performance than Aurora.
* C. Use Amazon RDS deployed in a Multi-AZ cluster deployment to create an Oracle database Di-rect the reporting functions to use the reader instance in the cluster deployment. This solution will not meet the requirement of improving application performance, as Oracle may not provide better performance than Aurora. It also involves using a cluster deployment, which is only supported for Aurora, not for Oracle. Reference URL: https://aws.amazon.com/rds/aurora/

NEW QUESTION 26
- (Topic 4)
A company runs a highly available SFTP service. The SFTP service uses two Amazon EC2
Linux instances that run with elastic IP addresses to accept traffic from trusted IP sources on the internet. The SFTP service is backed by shared storage that is attached to the instances. User accounts are created and managed as Linux users in the SFTP servers.
The company wants a serverless option that provides high IOPS performance and highly configurable security. The company also wants to maintain control over user permissions.
Which solution will meet these requirements?

A. Create an encrypted Amazon Elastic Block Store (Amazon EBS) volum

B. Create an AWS Transfer Family SFTP service with a public endpoint that allows only trusted IP addresse
C. Attach the EBS volume to the SFTP service endpoin
D. Grant users access to the SFTP service.
E. Create an encrypted Amazon Elastic File System (Amazon EFS) volum
F. Create an AWS Transfer Family SFTP service with elastic IP addresses and a VPC endpoint that has internet-facing acces
G. Attach a security group to the endpoint that allows only trusted IP addresse
H. Attach the EFS volume to the SFTP service endpoin
I. Grant users access to the SFTP service.
J. Create an Amazon S3 bucket with default encryption enable
K. Create an AWS Transfer Family SFTP service with a public endpoint that allows only trusted IP addresse
L. Attach the S3 bucket to the SFTP service endpoin
M. Grant users access to the SFTP service.
N. Create an Amazon S3 bucket with default encryption enable
O. Create an AWS Transfer Family SFTP service with a VPC endpoint that has internal access in a private subne
P. Attach a security group that allows only trusted IP addresse
Q. Attach the S3 bucket to the SFTP service endpoin
R. Grant users access to the SFTP service.

Answer: C

Explanation:
AWS Transfer Family is a secure transfer service that enables you to transfer files into and out of AWS storage services using SFTP, FTPS, FTP, and AS2 protocols. You can use AWS Transfer Family to create an SFTP-enabled server with a public endpoint that allows only trusted IP addresses. You can also attach an Amazon S3 bucket with default encryption enabled to the SFTP service endpoint, which will provide high IOPS performance and highly configurable security for your data at rest. You can also maintain control over user permissions by granting users access to the SFTP service using IAM roles or service-managed identities. References: https://docs.aws.amazon.com/transfer/latest/userguide/what-is-aws-transfer-family.html
https://docs.aws.amazon.com/transfer/latest/userguide/create-server-s3.html

NEW QUESTION 27
- (Topic 4)
A solutions architect is implementing a document review application using an Amazon S3 bucket for storage. The solution must prevent accidental deletion of the documents and ensure that all versions of the documents are available. Users must be able to download, modify, and upload documents.
Which combination of actions should be taken to meet these requirements? (Choose two.)

A. Enable a read-only bucket ACL.
B. Enable versioning on the bucket.
C. Attach an IAM policy to the bucket.
D. Enable MFA Delete on the bucket.
E. Encrypt the bucket using AWS KMS.

Answer: BD

Explanation:
Versioning is a feature of Amazon S3 that allows users to keep multiple versions of the same object in a bucket. It can help prevent accidental deletion of the documents and ensure that all versions of the documents are available1. MFA Delete is a feature of Amazon S3 that adds an extra layer of security by requiring two forms of authentication to delete a version or change the versioning state of a bucket. It can help prevent unauthorized or accidental deletion of the documents2. By enabling both versioning and MFA Delete on the bucket, the solution can meet the requirements.
* A. Enable a read-only bucket ACL. This solution will not meet the requirement of allowing
users to download, modify, and upload documents, as a read-only bucket ACL will prevent write access to the bucket3.
* C. Attach an IAM policy to the bucket. This solution will not meet the requirement of preventing accidental deletion of the documents and ensuring that all versions of the documents are available, as an IAM policy is used to grant or deny permissions to users or roles, not to enable versioning or MFA Delete4.
* E. Encrypt the bucket using AWS KMS. This solution will not meet the requirement of preventing accidental deletion of the documents and ensuring that all versions of the documents are available, as encrypting the bucket using AWS KMS is a method of protecting data at rest, not enabling versioning or MFA Delete.
Reference URL: https://docs.aws.amazon.com/AmazonS3/latest/userguide/Versioning.html

NEW QUESTION 31
- (Topic 4)
A company maintains about 300 TB in Amazon S3 Standard storage month after month The S3 objects are each typically around 50 GB in size and are frequently replaced with multipart uploads by their global application The number and size of S3 objects remain constant but the company's S3 storage costs are increasing each month.
How should a solutions architect reduce costs in this situation?

A. Switch from multipart uploads to Amazon S3 Transfer Acceleration.
B. Enable an S3 Lifecycle policy that deletes incomplete multipart uploads.
C. Configure S3 inventory to prevent objects from being archived too quickly.
D. Configure Amazon CloudFront to reduce the number of objects stored in Amazon S3.

Answer: B

Explanation:
This option is the most cost-effective way to reduce the S3 storage costs in this situation. Incomplete multipart uploads are parts of objects that are not completed or aborted by the application. They consume storage space and incur charges until they are deleted. By enabling an S3 Lifecycle policy that deletes incomplete multipart uploads, you can automatically remove them after a specified period of time (such as one day) and free up the storage space. This will reduce the S3 storage costs and also improve the performance of the application by avoiding unnecessary retries or errors.
Option A is not correct because switching from multipart uploads to Amazon S3 Transfer Acceleration will not reduce the S3 storage costs. Amazon S3 Transfer Acceleration is a feature that enables faster data transfers to and from S3 by using the AWS edge network. It is useful for improving the upload speed of large objects over long distances, but it does not affect the storage space or charges. In fact, it may increase the costs by adding a data transfer fee for using the feature.
Option C is not correct because configuring S3 inventory to prevent objects from being archived too quickly will not reduce the S3 storage costs. Amazon S3 Inventory is a feature that provides a report of the objects and their metadata in an S3 bucket. It is useful for managing and auditing the S3 objects, but it does not affect the storage space or charges. In fact, it may increase the costs by generating additional S3 objects for the inventory reports.

Option D is not correct because configuring Amazon CloudFront to reduce the number of objects stored in Amazon S3 will not reduce the S3 storage costs. Amazon CloudFront is a content delivery network (CDN) that distributes the S3 objects to edge locations for faster and lower latency access. It is useful for improving the download speed and availability of the S3 objects, but it does not affect the storage space or charges. In fact, it may increase the costs by adding a data transfer fee for using the service. References:
? Managing your storage lifecycle
? Using multipart upload
? Amazon S3 Transfer Acceleration
? Amazon S3 Inventory
? What Is Amazon CloudFront?

NEW QUESTION 33
- (Topic 4)
A company is reviewing a recent migration of a three-tier application to a VPC. The security team discovers that the principle of least privilege is not being applied to Amazon EC2 security group ingress and egress rules between the application tiers.
What should a solutions architect do to correct this issue?

A. Create security group rules using the instance ID as the source or destination.
B. Create security group rules using the security group ID as the source or destination.
C. Create security group rules using the VPC CIDR blocks as the source or destination.
D. Create security group rules using the subnet CIDR blocks as the source or destination.

Answer: B

Explanation:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/security-group- rules.html

NEW QUESTION 36
- (Topic 4)
A company wants to migrate an on-premises legacy application to AWS. The application ingests customer order files from an on-premises enterprise resource planning (ERP) system. The application then uploads the files to an SFTP server. The application uses a scheduled job that checks for order files every hour.
The company already has an AWS account that has connectivity to the on-premises network. The new application on AWS must support integration with the existing ERP system. The new application must be secure and resilient and must use the SFTP protocol to process orders from the ERP system immediately.
Which solution will meet these requirements?

A. Create an AWS Transfer Family SFTP internet-facing server in two Availability Zone
B. Use Amazon S3 storag
C. Create an AWS Lambda function to process order file
D. Use S3 Event Notifications to send s3: ObjectCreated: * events to the Lambda function.
E. Create an AWS Transfer Family SFTP internet-facing server in one Availability Zon
F. Use Amazon Elastic File System (Amazon EFS) storag
G. Create an AWS Lambda function to process order file
H. Use a Transfer Family managed workflow to invoke the Lambda function.
I. Create an AWS Transfer Family SFTP internal server in two Availability Zone
J. Use Amazon Elastic File System (Amazon EFS) storag
K. Create an AWS Step Functions state machine to process order file
L. Use Amazon EventBridge Scheduler to invoke the state machine to periodically check Amazon EFS for order files.
M. Create an AWS Transfer Family SFTP internal server in two Availability Zone
N. Use Amazon S3 storag
O. Create an AWS Lambda function to process order file
P. Use a Transfer Family managed workflow to invoke the Lambda function.

Answer: D

Explanation:
This solution meets the requirements because it uses the following components and features:
? AWS Transfer Family SFTP internal server: This allows the application to securely
transfer order files from the on-premises ERP system to AWS using the SFTP protocol over a private connection. The internal server is deployed in two Availability Zones for high availability and fault tolerance.
? Amazon S3 storage: This provides scalable, durable, and cost-effective object
storage for the order files. Amazon S3 also supports encryption at rest and in transit, as well as lifecycle policies and versioning for data protection and compliance.
? AWS Lambda function: This enables the application to process the order files in a
serverless manner, without provisioning or managing servers. The Lambda function can perform any custom logic or transformation on the order files, such as validating, parsing, or enriching the data.
? Transfer Family managed workflow: This simplifies the orchestration of the file
processing tasks by triggering the Lambda function as soon as a file is uploaded to the SFTP server. The managed workflow also provides error handling, retry policies, and logging capabilities.

NEW QUESTION 38
- (Topic 4)
A company is building an Amazon Elastic Kubernetes Service (Amazon EKS) cluster for its workloads. All secrets that are stored in Amazon EKS must be encrypted in the Kubernetes etcd key-value store.
Which solution will meet these requirements?

A. Create a new AWS Key Management Service (AWS KMS) key Use AWS Secrets Manager to manage rotate, and store all secrets in Amazon EKS.
B. Create a new AWS Key Management Service (AWS KMS) key Enable Amazon EKS KMS secrets encryption on the Amazon EKS cluster.
C. Create the Amazon EKS cluster with default options Use the Amazon Elastic Block Store (Amazon EBS) Container Storage Interface (CSI) driver as an add-on.
D. Create a new AWS Key Management Service (AWS KMS) key with the ahas/aws/ebs alias Enable default Amazon Elastic Block Store (Amazon EBS) volume encryption for the account.

Answer: B

Explanation:
This option is the most secure and simple way to encrypt the secrets that are stored in Amazon EKS. AWS Key Management Service (AWS KMS) is a service that allows you to create and manage encryption keys that can be used to encrypt your data. Amazon EKS KMS secrets encryption is a feature that enables you to use a KMS key to encrypt the secrets that are stored in the Kubernetes etcd key-value store. This provides an additional layer of protection for your sensitive data, such as passwords, tokens, and keys. You can create a new KMS key or use an existing one, and then enable the Amazon EKS KMS secrets encryption on the Amazon EKS cluster. You can also use IAM policies to control who can access or use the KMS key.
Option A is not correct because using AWS Secrets Manager to manage, rotate, and store all secrets in Amazon EKS is not necessary or efficient. AWS Secrets Manager is a service that helps you securely store, retrieve, and rotate your secrets, such as database credentials, API keys, and passwords. You can use it to manage secrets that are used by your applications or services outside of Amazon EKS, but it is not designed to encrypt the secrets that are stored in the Kubernetes etcd key-value store. Moreover, using AWS Secrets Manager would incur additional costs and complexity, and it would not leverage the native Kubernetes secrets management capabilities.
Option C is not correct because using the Amazon EBS Container Storage Interface (CSI) driver as an add-on does not encrypt the secrets that are stored in Amazon EKS. The Amazon EBS CSI driver is a plugin that allows you to use Amazon EBS volumes as persistent storage for your Kubernetes pods. It is useful for providing durable and scalable storage for your applications, but it does not affect the encryption of the secrets that are stored in the Kubernetes etcd key-value store. Moreover, using the Amazon EBS CSI driver would require additional configuration and resources, and it would not provide the same level of security as using a KMS key.
Option D is not correct because creating a new AWS KMS key with the alias aws/ebs and enabling default Amazon EBS volume encryption for the account does not encrypt the secrets that are stored in Amazon EKS. The alias aws/ebs is a reserved alias that is used by AWS to create a default KMS key for your account. This key is used to encrypt the Amazon EBS volumes that are created in your account, unless you specify a different KMS key. Enabling default Amazon EBS volume encryption for the account is a setting that ensures that all new Amazon EBS volumes are encrypted by default. However, these features do not affect the encryption of the secrets that are stored in the Kubernetes etcd key-value store. Moreover, using the default KMS key or the default encryption setting would not provide the same level of control and security as using a custom KMS key and enabling the Amazon EKS KMS secrets encryption feature. References:
? Encrypting secrets used in Amazon EKS
? What Is AWS Key Management Service?
? What Is AWS Secrets Manager?
? Amazon EBS CSI driver
? Encryption at rest

NEW QUESTION 39
- (Topic 4)
A company stores multiple Amazon Machine Images (AMIs) in an AWS account to launch its Amazon EC2 instances. The AMIs contain critical data and configurations that are necessary for the company's operations. The company wants to implement a solution that will recover accidentally deleted AMIs quickly and efficiently.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create Amazon Elastic Block Store (Amazon EBS) snapshots of the AMI
B. Store the snapshots in a separate AWS account.
C. Copy all AMIs to another AWS account periodically.
D. Create a retention rule in Recycle Bin.
E. Upload the AMIs to an Amazon S3 bucket that has Cross-Region Replication.

Answer: C

Explanation:
Recycle Bin is a data recovery feature that enables you to restore accidentally deleted Amazon EBS snapshots and EBS-backed AMIs. When using Recycle Bin, if your resources are deleted, they are retained in the Recycle Bin for a time period that you specify before being permanently deleted. You can restore a resource from the Recycle Bin at any time before its retention period expires. This solution has the least operational overhead, as you do not need to create, copy, or upload any additional resources. You can also manage tags and permissions for AMIs in the Recycle Bin. AMIs in the Recycle Bin do not incur any additional charges.
References:
? Recover AMIs from the Recycle Bin
? Recover an accidentally deleted Linux AMI

NEW QUESTION 41
- (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 is concerned that two NAT instances in use will no longer be able to support the traffic needed for the company’s application. A solutions architect wants to implement a solution that is highly available, fault tolerant, and automatically scalable.

What should the solutions architect recommend?

A. Remove the two NAT instances and replace them with two NAT gateways in the same Availability Zone.
B. Use Auto Scaling groups with Network Load Balancers for the NAT instances in differentAvailability Zones.
C. Remove the two NAT instances and replace them with two NAT gateways in different Availability Zones.
D. Replace the two NAT instances with Spot Instances in different Availability Zones and deploy a Network Load Balancer.

Answer: C

Explanation:
If you have resources in multiple Availability Zones and they share one NAT gateway, and if the NAT gateway's Availability Zone is down, resources in the other Availability Zones lose internet access. To create an Availability Zone-independent architecture, create a NAT gateway in each Availability Zone and configure your routing to ensure that resources use the NAT gateway in the same Availability Zone. https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html#nat-gateway- basics

NEW QUESTION 47
- (Topic 4)
A retail company uses a regional Amazon API Gateway API for its public REST APIs. The API Gateway endpoint is a custom domain name that points to an Amazon Route 53 alias record. A solutions architect needs to create a solution that has minimal effects on customers and minimal data loss to release the new version of APIs.
Which solution will meet these requirements?

A. Create a canary release deployment stage for API Gatewa
B. Deploy the latest API versio
C. Point an appropriate percentage of traffic to the canary stag
D. After API verification, promote the canary stage to the production stage.
E. Create a new API Gateway endpoint with a new version of the API in OpenAPI YAMLfile forma
F. Use the import-to-update operation in merge mode into the API in API Gatewa
G. Deploy the new version of the API to the production stage.
H. Create a new API Gateway endpoint with a new version of the API in OpenAPI JSON file forma
I. Use the import-to-update operation in overwrite mode into the API in API Gatewa
J. Deploy the new version of the API to the production stage.
K. Create a new API Gateway endpoint with new versions of the API definition
L. Create a custom domain name for the new API Gateway AP
M. Point the Route 53 alias record to the new API Gateway API custom domain name.

Answer: A

Explanation:
This answer is correct because it meets the requirements of releasing the new version of APIs with minimal effects on customers and minimal data loss. A canary release deployment is a software development strategy in which a new version of an API is deployed for testing purposes, and the base version remains deployed as a production release for normal operations on the same stage. In a canary release deployment, total API traffic is separated at random into a production release and a canary release with a pre- configured ratio. Typically, the canary release receives a small percentage of API traffic and the production release takes up the rest. The updated API features are only visible to API traffic through the canary. You can adjust the canary traffic percentage to optimize test coverage or performance. By keeping canary traffic small and the selection random, most users are not adversely affected at any time by potential bugs in the new version, and no single user is adversely affected all the time. After the test metrics pass your requirements, you can promote the canary release to the production release and disable the canary from the deployment. This makes the new features available in the production stage. References:
? https://docs.aws.amazon.com/apigateway/latest/developerguide/canary-release.html

NEW QUESTION 51
- (Topic 4)
A company built an application with Docker containers and needs to run the application in the AWS Cloud The company wants to use a managed sen/ice to host the application
The solution must scale in and out appropriately according to demand on the individual container services The solution also must not result in additional operational overhead or infrastructure to manage
Which solutions will meet these requirements? (Select TWO)

A. Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate.
B. Use Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate.
C. Provision an Amazon API Gateway API Connect the API to AWS Lambda to run thecontainers.
D. Use Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 worker nodes.
E. Use Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 worker nodes.

Answer: AB

Explanation:
These options are the best solutions because they allow the company to run the application with Docker containers in the AWS Cloud using a managed service that scales automatically and does not require any infrastructure to manage. By using AWS Fargate, the company can launch and run containers without having to provision, configure, or scale clusters of EC2 instances. Fargate allocates the right amount of compute resources for each container and scales them up or down as needed. By using Amazon ECS or Amazon EKS, the company can choose the container orchestration platform that suits its needs. Amazon ECS is a fully managed service that integrates with other AWS services and simplifies the deployment and management of containers. Amazon EKS is a managed service that runs Kubernetes on AWS and provides compatibility with existing Kubernetes tools and plugins.
* C. Provision an Amazon API Gateway API Connect the API to AWS Lambda to run the containers. This option is not feasible because AWS Lambda does not support running Docker containers directly. Lambda functions are executed in a sandboxed environment that is isolated from other functions and resources. To run Docker containers on Lambda, the company would need to use a custom runtime or a wrapper library that emulates the Docker API, which can introduce additional complexity and overhead.
* D. Use Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 worker nodes. This option is not optimal because it requires the company to manage the EC2 instances that host 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.
* E. Use Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 worker nodes. This option is not ideal because it requires the company to manage the EC2 instances that host 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.
References:
? 1 AWS Fargate - Amazon Web Services
? 2 Amazon Elastic Container Service - Amazon Web Services
? 3 Amazon Elastic Kubernetes Service - Amazon Web Services
? 4 AWS Lambda FAQs - Amazon Web Services

NEW QUESTION 53
- (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 56
- (Topic 4)
A company needs to integrate with a third-party data feed. The data feed sends a webhook to notify an external service when new data is ready for consumption A developer wrote an AWS Lambfe function to retrieve data when the company receives a webhook callback The developer must make the Lambda function available for the third party to call.
Which solution will meet these requirements with the MOST operational efficiency?

A. Create a function URL for the Lambda functio
B. Provide the Lambda function URL to the third party for the webhook.
C. Deploy an Application Load Balancer (ALB) in front of the Lambda functio
D. Provide the ALB URL to the third party for the webhook
E. Create an Amazon Simple Notification Service (Amazon SNS) topi
F. Attach the topic to the Lambda functio
G. Provide the public hostname of the SNS topic to the third party for the webhook.
H. Create an Amazon Simple Queue Service (Amazon SQS) queu
I. Attach the queue to the Lambda functio
J. Provide the public hostname of the SQS queue to the third party forthe webhook.

Answer: A

Explanation:
A function URL is a unique identifier for a Lambda function that can be used to invoke the function over HTTPS. It is composed of the API endpoint of the AWS Region where the function is deployed, and the name or ARN of the function1. By creating a function URL for the Lambda function, the solution can make the Lambda function available for the third party to call with the most operational efficiency.
* B. Deploy an Application Load Balancer (ALB) in front of the Lambda function. Provide the ALB URL to the third party for the webhook. This solution will not meet the requirement of the most operational efficiency, as it involves creating and managing an additional resource (ALB) that is not necessary for invoking a Lambda function over HTTPS2.
* C. Create an Amazon Simple Notification Service (Amazon SNS) topic. Attach the topic to the Lambda function. Provide the public hostname of the SNS topic to the third party for the webhook. This solution will not work, as Amazon SNS topics do not have public hostnames that can be used as webhooks. SNS topics are used to publish messages to subscribers, not to receive messages from external sources3.
* D. Create an Amazon Simple Queue Service (Amazon SQS) queue. Attach the queue to the Lamb-da function. Provide the public hostname of the SQS queue to the third party for the webhook. This solution will not work, as Amazon SQS queues do not have public hostnames that can be used as webhooks. SQS queues are used to send, store, and receive messages between AWS services, not to receive messages from external sources. Reference URL: https://docs.aws.amazon.com/lambda/latest/dg/lambda-api-permissions- ref.html

NEW QUESTION 60
- (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 62
- (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 65
- (Topic 4)
A company sends AWS CloudTrail logs from multiple AWS accounts to an Amazon S3 bucket in a centralized account. The company must keep the CloudTrail logs. The company must also be able to query the CloudTrail logs at any time
Which solution will meet these requirements?

A. Use the CloudTrail event history in the centralized account to create an Amazon Athena tabl
B. Query the CloudTrail logs from Athena.
C. Configure an Amazon Neptune instance to manage the CloudTrail log
D. Query the CloudTrail logs from Neptune.
E. Configure CloudTrail to send the logs to an Amazon DynamoDB tabl
F. Create a dashboard in Amazon QulCkSight to query the logs in the table.
G. use Amazon Athena to create an Athena noteboo
H. Configure CloudTrail to send the logs to the noteboo
I. Run queries from Athena.

Answer: A

Explanation:
it allows the company to keep the CloudTrail logs and query them at any time. By using the CloudTrail event history in the centralized account, the company can view, filter, and download recent API activity across multiple AWS accounts. By creating an Amazon Athena table from the CloudTrail event history, the company can use a serverless interactive query service that makes it easy to analyze data in S3 using standard SQL. By querying the CloudTrail logs from Athena, the company can gain insights into user activity and resource changes. References:
? Viewing Events with CloudTrail Event History
? Querying AWS CloudTrail Logs
? Amazon Athena

NEW QUESTION 67
- (Topic 4)
A company is deploying an application that processes streaming data in near-real time The company plans to use Amazon EC2 instances for the workload The network architecture must be configurable to provide the lowest possible latency between nodes
Which combination of network solutions will meet these requirements? (Select TWO)

A. Enable and configure enhanced networking on each EC2 instance
B. Group the EC2 instances in separate accounts
C. Run the EC2 instances in a cluster placement group
D. Attach multiple elastic network interfaces to each EC2 instance
E. Use Amazon Elastic Block Store (Amazon EBS) optimized instance types.

Answer: AC

Explanation:
These options are the most suitable ways to configure the network architecture to provide the lowest possible latency between nodes. Option A enables and configures enhanced networking on each EC2 instance, which is a feature that improves the network performance of the instance by providing higher bandwidth, lower latency, and lower jitter. Enhanced networking uses single root I/O virtualization (SR-IOV) or Elastic Fabric Adapter (EFA) to provide direct access to the network hardware. You can enable and configure enhanced networking by choosing a supported instance type and a compatible operating system, and installing the required drivers. Option C runs the EC2 instances in a cluster placement group, which is a logical grouping of instances within a single Availability Zone that are placed close together on the same underlying hardware. Cluster placement groups provide the lowest network latency and the highest network throughput among the placement group options. You can run the EC2 instances in a cluster placement group by creating a placement group and launching the instances into it. Option B is not suitable because grouping the EC2 instances in separate accounts does not provide the lowest possible latency between nodes. Separate accounts are used to isolate and organize resources for different purposes, such as security, billing, or compliance. However, they do not affect the network performance or proximity of the instances. Moreover, grouping the EC2 instances in separate accounts would incur additional costs and complexity, and it would require setting up cross-account networking and permissions.
Option D is not suitable because attaching multiple elastic network interfaces to each EC2 instance does not provide the lowest possible latency between nodes. Elastic network interfaces are virtual network interfaces that can be attached to EC2 instances to provide additional network capabilities, such as multiple IP addresses, multiple subnets, or enhanced security. However, they do not affect the network performance or proximity of the instances. Moreover, attaching multiple elastic network interfaces to each EC2 instance would consume additional resources and limit the instance type choices.
Option E is not suitable because using Amazon EBS optimized instance types does not provide the lowest possible latency between nodes. Amazon EBS optimized instance types are instances that provide dedicated bandwidth for Amazon EBS volumes, which are block storage volumes that can be attached to EC2 instances. EBS optimized instance types improve the performance and consistency of the EBS volumes, but they do not affect the network performance or proximity of the instances. Moreover, using EBS optimized instance types would incur additional costs and may not be necessary for the streaming data workload.
References:
? Enhanced networking on Linux
? Placement groups
? Elastic network interfaces
? Amazon EBS-optimized instances

NEW QUESTION 70
- (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 71
- (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 76
- (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 81
- (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 82
- (Topic 4)
A company needs to store data from its healthcare application. The application's data frequently changes. A new regulation requires audit z access at all levels of the stored data.
The company hosts the application on an on-premises infrastructure that is running out of storage capacity. A solutions architect must securely migrate the existing data to AWS while satisfying the new regulation.
Which solution will meet these requirements?

A. Use AWS DataSync to move the existing data to Amazon S3. Use AWS CloudTrail to log data events.
B. Use AWS Snowcone to move the existing data to Amazon $3. Use AWS CloudTrail to log management events.
C. Use Amazon S3 Transfer Acceleration to move the existing data to Amazon S3. Use AWS CloudTrail to log data events.
D. Use AWS Storage Gateway to move the existing data to Amazon S3. Use AWS CloudTrail to log management events.

Answer: A

Explanation:
This answer is correct because it meets the requirements of securely migrating the existing data to AWS and satisfying the new regulation. AWS DataSync is a service that makes it easy to move large amounts of data online between on-premises storage and Amazon S3. DataSync automatically encrypts data in transit and verifies data integrity during transfer. AWS CloudTrail is a service that records AWS API calls for your account and delivers log files to Amazon S3. CloudTrail can log data events, which show the resource operations performed on or within a resource in your AWS account, such as S3 object-level API activity. By using CloudTrail to log data events, you can audit access at all levels of the stored data.
References:
? https://docs.aws.amazon.com/datasync/latest/userguide/what-is-datasync.html
? https://docs.aws.amazon.com/awscloudtrail/latest/userguide/logging-data-events- with-cloudtrail.html

NEW QUESTION 87
- (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 90
- (Topic 4)
A company is running a microservices application on Amazon EC2 instances. The company wants to migrate the application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster for scalability. The company must configure the Amazon EKS control plane with endpoint private access set to true and endpoint public access set to false to maintain security compliance The company must also put the data plane in private subnets. However, the company has received error notifications because the node cannot join the cluster.
Which solution will allow the node to join the cluster?

A. Grant the required permission in AWS Identity and Access Management (1AM) to the AmazonEKSNodeRole 1AM role.
B. Create interface VPC endpoints to allow nodes to access the control plane.
C. Recreate nodes in the public subnet Restrict security groups for EC2 nodes
D. Allow outbound traffic in the security group of the nodes.

Answer: B

Explanation:
Kubernetes API requests within your cluster's VPC (such as node to control plane communication) use the private VPC endpoint.
https://docs.aws.amazon.com/eks/latest/userguide/cluster-endpoint.html

NEW QUESTION 91
- (Topic 4)
A company has multiple Windows file servers on premises. The company wants to migrate and consolidate its files into an Amazon FSx for Windows File Server file system. File permissions must be preserved to ensure that access rights do not change.
Which solutions will meet these requirements? (Select TWO.)

A. Deploy AWS DataSync agents on premise
B. Schedule DataSync tasks to transfer the data to the FSx for Windows File Server file system.
C. Copy the shares on each file server into Amazon S3 buckets by using the AWS CLI Schedule AWS DataSync tasks to transfer the data to the FSx for Windows File Server file system.
D. Remove the drives from each file server Ship the drives to AWS for import into Amazon S3. Schedule AWS DataSync tasks to transfer the data to the FSx for Windows File Server file system
E. Order an AWS Snowcone devic
F. Connect the device to the on-premises networ
G. Launch AWS DataSync agents on the devic
H. Schedule DataSync tasks to transfer the data to the FSx for Windows File Server file system,
I. Order an AWS Snowball Edge Storage Optimized devic
J. Connect the device to the on- premises networ
K. Copy data to the device by using the AWS CL
L. Ship the device back to AWS for import into Amazon S3. Schedule AWS DataSync tasks to transfer the data to the FSx for Windows File Server file system.

Answer: AD

Explanation:

A This option involves deploying DataSync agents on your on-premises file servers and using DataSync to transfer the data directly to the FSx for Windows File Server. DataSync ensures that file permissions are preserved during the migration process. D This option involves using an AWS Snowcone device, a portable data transfer device. You would connect the Snowcone device to your on-premises network, launch DataSync agents on the device, and schedule DataSync tasks to transfer the data to FSx for Windows File Server. DataSync handles the migration process while preserving file permissions.

NEW QUESTION 92
- (Topic 4)
A company has a large workload that runs every Friday evening. The workload runs on Amazon EC2 instances that are in two Availability Zones in the us-east-1 Region. Normally, the company must run no more than two instances at all times. However, the company wants to scale up to six instances each Friday to handle a regularly repeating increased workload.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create a reminder in Amazon EventBridge to scale the instances.
B. Create an Auto Scaling group that has a scheduled action.
C. Create an Auto Scaling group that uses manual scaling.
D. Create an Auto Scaling group that uses automatic scaling.

Answer: B

Explanation:
An Auto Scaling group is a collection of EC2 instances that share similar characteristics and can be scaled in or out automatically based on demand. An Auto Scaling group can have a scheduled action, which is a configuration that tells the group to scale to a specific size at a specific time. This way, the company can scale up to six instances each Friday evening to handle the increased workload, and scale down to two instances at other times to save costs. This solution meets the requirements with the least operational overhead, as it does not require manual intervention or custom scripts. References:
? 1 explains how to create a scheduled action for an Auto Scaling group.
? 2 describes the concept and benefits of an Auto Scaling group.

NEW QUESTION 97
- (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 98
- (Topic 4)
The customers of a finance company request appointments with financial advisors by sending text messages. A web application that runs on Amazon EC2 instances accepts the appointment requests. The text messages are published to an Amazon Simple Queue Service (Amazon SQS) queue through the web application. Another application that runs on EC2 instances then sends meeting invitations and meeting confirmation email messages to the customers. After successful scheduling, this application stores the meeting information in an Amazon DynamoDB database.
As the company expands, customers report that their meeting invitations are taking longer to arrive.
What should a solutions architect recommend to resolve this issue?

A. Add a DynamoDB Accelerator (DAX) cluster in front of the DynamoDB database.
B. Add an Amazon API Gateway API in front of the web application that accepts the appointment requests.

C. Add an Amazon CloudFront distributio
D. Set the origin as the web application that accepts the appointment requests.
E. Add an Auto Scaling group for the application that sends meeting invitation
F. Configure the Auto Scaling group to scale based on the depth of the SQS queue.

Answer: D

Explanation:
To resolve the issue of longer delivery times for meeting invitations, the solutions architect can recommend adding an Auto Scaling group for the application that sends meeting invitations and configuring the Auto Scaling group to scale based on the depth of the SQS queue. This will allow the application to scale up as the number of appointment requests increases, improving the performance and delivery times of the meeting invitations.

NEW QUESTION 99
- (Topic 4)
An loT company is releasing a mattress that has sensors to collect data about a user's sleep. The sensors will send data to an Amazon S3 bucket. The sensors collect approximately 2 MB of data every night for each mattress. The company must process and summarize the data for each mattress. The results need to be available as soon as possible Data processing will require 1 GB of memory and will finish within 30 seconds.
Which solution will meet these requirements MOST cost-effectively?

A. Use AWS Glue with a Scalajob.
B. Use Amazon EMR with an Apache Spark script.
C. Use AWS Lambda with a Python script.
D. Use AWS Glue with a PySpark job.

Answer: C

Explanation:
AWS Lambda charges you based on the number of invocations and the execution time of your function. Since the data processing job is relatively small (2 MB of data), Lambda is a cost-effective choice. You only pay for the actual usage without the need to provision and maintain infrastructure.

NEW QUESTION 101
- (Topic 4)
A company uses AWS Organizations. The company wants to operate some of its AWS accounts with different budgets. The company wants to receive alerts and automatically prevent provisioning of additional resources on AWS accounts when the allocated budget threshold is met during a specific period.
Which combination of solutions will meet these requirements? (Select THREE.)

A. Use AWS Budgets to create a budge
B. Set the budget amount under the Cost and Usage Reports section of the required AWS accounts.
C. Use AWS Budgets to create a budge
D. Set the budget amount under the Billing dashboards of the required AWS accounts.
E. Create an 1AM user for AWS Budgets to run budget actions with the required permissions.
F. Create an 1AM role for AWS Budgets to run budget actions with the required permissions.
G. Add an alert to notify the company when each account meets its budget threshol
H. Add a budget action that selects the 1AM identity created with the appropriate config rule to prevent provisioning of additional resources.
I. Add an alert to notify the company when each account meets its budget threshol
J. Add a budget action that selects the 1AM identity created with the appropriate service control policy (SCP) to prevent provisioning of additional resources.

Answer: BDF

Explanation:
To use AWS Budgets to create and manage budgets for different AWS accounts, the company needs to do the following steps:
? Use AWS Budgets to create a budget for each AWS account that needs a different
budget amount. The budget can be based on cost or usage metrics, and can have different time periods, filters, and thresholds. The company can set the budget amount under the Billing dashboards of the required AWS accounts1.
? Create an IAM role for AWS Budgets to run budget actions with the required
permissions. A budget action is a response that AWS Budgets initiates when a
budget exceeds a specified threshold. The IAM role allows AWS Budgets to perform actions on behalf of the company, such as applying an IAM policy or a service control policy (SCP) to restrict the provisioning of additional resources2.
? Add an alert to notify the company when each account meets its budget threshold.
The alert can be sent via email or Amazon SNS. The company can also add a budget action that selects the IAM role created and the appropriate SCP to prevent provisioning of additional resources. An SCP is a type of policy that can be applied to an AWS account or an organizational unit (OU) within AWS Organizations. An SCP can limit the actions that users and roles can perform in the account or OU3.
References:
? 4: https://aws.amazon.com/budgets/
? 1: https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/budgets- create.html
? 2: https://docs.aws.amazon.com/cost-management/latest/userguide/budgets- controls.html
? 3:
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policie s_scps.html

NEW QUESTION 104
- (Topic 4)
A company hosts a website on Amazon EC2 instances behind an Application Load Balancer (ALB) The website serves static content Website traffic is increasing and the company is concerned about a potential increase in cost.
What should a solutions architect do to reduce the cost of the website?

A. Create an Amazon CloudFront distribution to cache static files at edge locations.
B. Create an Amazon ElastiCache cluster Connect the ALB to the ElastiCache cluster to serve cached files.
C. Create an AWS WAF web ACL and associate it with the AL
D. Add a rule to the web ACL to cache static files.
E. Create a second ALB in an alternative AWS Region Route user traffic to the closest Region to minimize data transfer costs

Answer: A

Explanation:
Amazon CloudFront is a content delivery network (CDN) that can improve the performance and reduce the cost of serving static content from a website. CloudFront
can cache static files at edge locations closer to the users, reducing the latency and data transfer costs. CloudFront can also integrate with Amazon S3 as the origin for the static content, eliminating the need for EC2 instances to host the website. CloudFront meets all the requirements of the question, while the other options do not. References:
? https://aws.amazon.com/blogs/architecture/architecting-a-low-cost-web-content-publishing-system/
? https://nodeployfriday.com/posts/static-website-hosting/
? https://aws.amazon.com/cloudfront/

NEW QUESTION 106
- (Topic 4)
A solutions architect must provide an automated solution for a company's compliance policy that states security groups cannot include a rule that allows SSH from 0.0.0.0/0. The company needs to be notified if there is any breach in the policy. A solution is needed as soon as possible.
What should the solutions architect do to meet these requirements with the LEAST operational overhead?

A. Write an AWS Lambda script that monitors security groups for SSH being open to 0.0.0.0/0 addresses and creates a notification every time it finds one.
B. Enable the restricted-ssh AWS Config managed rule and generate an Amazon Simple Notification Service (Amazon SNS) notification when a noncompliant rule is created.
C. Create an 1AM role with permissions to globally open security groups and network ACL
D. Create an Amazon Simple Notification Service (Amazon SNS) topic to generate a notification every time the role is assumed by a user.
E. Configure a service control policy (SCP) that prevents non-administrative users from creating or editing security group
F. Create a notification in the ticketing system when a user requests a rule that needs administrator permissions.

Answer: B

Explanation:
The most suitable solution for the company's compliance policy is to enable the restricted-ssh AWS Config managed rule and generate an Amazon Simple Notification Service (Amazon SNS) notification when a noncompliant rule is created. This solution has the least operational overhead because it uses a predefined rule that is already available in AWS Config, which is a service that enables users to assess, audit, and evaluate the configurations of their AWS resources. The restricted-ssh rule checks whether security groups that are in use have inbound rules that allow SSH from 0.0.0.0/0 addresses, and reports them as noncompliant1. Users can configure the rule to send notifications to an Amazon SNS topic when a noncompliant change occurs, and subscribe to the topic to receive alerts via email, SMS, or other methods2.
The other options are not correct because they either have more operational overhead or do not meet the requirements. Writing an AWS Lambda script that monitors security groups for SSH being open to 0.0.0.0/0 addresses and creates a notification every time it finds one is not correct because it requires custom code development and maintenance, which adds complexity and cost to the solution. Creating an IAM role with permissions to globally open security groups and network ACLs, and creating an Amazon SNS topic to generate a notification every time the role is assumed by a user is not correct because it does not prevent or detect the creation of noncompliant rules by other users or roles, and it does not address the existing rules that may violate the policy. Configuring a service control policy (SCP) that prevents non-administrative users from creating or editing security groups, and creating a notification in the ticketing system when a user requests a rule that needs administrator permissions is not correct because it does not provide an automated solution for the policy enforcement and notification, and it may limit the flexibility and productivity of the users.
References:
? restricted-ssh - AWS Config
? Getting Notifications When Your Resources Change - AWS Config

NEW QUESTION 107
- (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 108
- (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 113
- (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 114
- (Topic 4)
An ecommerce company runs applications in AWS accounts that are part of an organization in AWS Organizations The applications run on Amazon Aurora PostgreSQL databases across all the accounts The company needs to prevent malicious activity and must identify abnormal failed and incomplete login attempts to the databases
Which solution will meet these requirements in the MOST operationally efficient way?

A. Attach service control policies (SCPs) to the root of the organization to identify the failed login attempts
B. Enable the Amazon RDS Protection feature in Amazon GuardDuty for the member accounts of the organization
C. Publish the Aurora general logs to a log group in Amazon CloudWatch Logs Export the log data to a central Amazon S3 bucket
D. Publish all the Aurora PostgreSQL database events in AWS CloudTrail to a central Amazon S3 bucket

Answer: C

Explanation:
This option is the most operationally efficient way to meet the requirements because it allows the company to monitor and analyze the database login activity across all the accounts in the organization. By publishing the Aurora general logs to a log group in Amazon CloudWatch Logs, the company can enable the logging of the database connections, disconnections, and failed authentication attempts. By exporting the log data to a central Amazon S3 bucket, the company can store the log data in a durable and cost- effective way and use other AWS services or tools to perform further analysis or alerting on the log data. For example, the company can use Amazon Athena to query the log data in Amazon S3, or use Amazon SNS to send notifications based on the log data.
* A. Attach service control policies (SCPs) to the root of the organization to identify the failed login attempts. This option is not effective because SCPs are not designed to identify the failed login attempts, but to restrict the actions that the users and roles can perform in the member accounts of the organization. SCPs are applied to the AWS API calls, not to the database login attempts. Moreover, SCPs do not provide any logging or analysis capabilities for the database activity.
* B. Enable the Amazon RDS Protection feature in Amazon GuardDuty for the member accounts of the organization. This option is not optimal because the Amazon RDS Protection feature in Amazon GuardDuty is not available for Aurora PostgreSQL databases, but only for Amazon RDS for MySQL and Amazon RDS for MariaDB databases. Moreover, the Amazon RDS Protection feature does not monitor the database login attempts, but the network and API activity related to the RDS instances.
* D. Publish all the Aurora PostgreSQL database events in AWS CloudTrail to a central Amazon S3 bucket. This option is not sufficient because AWS CloudTrail does not capture the database login attempts, but only the AWS API calls made by or on behalf of the Aurora PostgreSQL database. For example, AWS CloudTrail can record the events such as creating, modifying, or deleting the database instances, clusters, or snapshots, but not the events such as connecting, disconnecting, or failing to authenticate to the database. References:
? 1 Working with Amazon Aurora PostgreSQL - Amazon Aurora
? 2 Working with log groups and log streams - Amazon CloudWatch Logs
? 3 Exporting Log Data to Amazon S3 - Amazon CloudWatch Logs

? [4] Amazon GuardDuty FAQs
? [5] Logging Amazon RDS API Calls with AWS CloudTrail - Amazon Relational Database Service

NEW QUESTION 115
- (Topic 4)
A company is storing 700 terabytes of data on a large network-attached storage (NAS) system in fts corporate data center. The company has a hybrid environment with a 10 Gbps AWS Direct Connect connection.
After an audit from a regulator, the company has 90 days to move the data to the cloud. The company needs to move the data efficiently and without disruption. The company still needs to be able to access and update the data during the transfer window.
Which solution will meet these requirements?

A. Create an AWS DataSync agent in the corporate data cente
B. Create a data transfer tas
C. Start the transfer to an Amazon S3 bucket.
D. Back up the data to AWS Snowball Edge Storage Optimized device
E. Ship the devices to an AWS data cente
F. Mount a target Amazon S3 bucket on the on-premises file system.
G. Use rsync to copy the data directly from local storage to a designated Amazon S3 bucket over the Direct Connect connection.
H. Back up the data on tape
I. Ship the tapes to an AWS data cente
J. Mount a target Amazon S3 bucket on the on-premises file system.

Answer: A

Explanation:
This answer is correct because it meets the requirements of moving the data efficiently and without disruption, and still being able to access and update the data during the transfer window. AWS DataSync is an online data movement and discovery service that simplifies and accelerates data migrations to AWS and helps you move data quickly and securely between on-premises storage, edge locations, other clouds, and AWS Storage. You can create an AWS DataSync agent in the corporate data center to connect your NAS system to AWS over the Direct Connect connection. You can create a data transfer task to specify the source location, destination location, and options for transferring the data. You can start the transfer to an Amazon S3 bucket and monitor the progress of the task. DataSync automatically encrypts data in transit and verifies data integrity during transfer. DataSync also supports incremental transfers, which means that only files that have changed since the last transfer are copied. This way, you can ensure that your data is synchronized between your NAS system and S3 bucket, and you can access and update the data during the transfer window.
References:
? https://docs.aws.amazon.com/datasync/latest/userguide/what-is-datasync.html
? https://docs.aws.amazon.com/datasync/latest/userguide/how-datasync-works.html

NEW QUESTION 116
- (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 117
- (Topic 4)
A company runs container applications by using Amazon Elastic Kubernetes Service (Amazon EKS). The company's workload is not consistent throughout the day The company wants Amazon EKS to scale in and out according to the workload.
Which combination of steps will meet these requirements with the LEAST operational overhead? {Select TWO.)

A. Use an AWS Lambda function to resize the EKS cluster
B. Use the Kubernetes Metrics Server to activate horizontal pod autoscaling.
C. Use the Kubernetes Cluster Autoscaler to manage the number of nodes in the cluster.
D. Use Amazon API Gateway and connect it to Amazon EKS
E. Use AWS App Mesh to observe network activity.

Answer: BC

Explanation:
https://docs.aws.amazon.com/eks/latest/userguide/horizontal-pod- autoscaler.html https://docs.aws.amazon.com/eks/latest/userguide/autoscaling.html
Horizontal pod autoscaling is a feature of Kubernetes that automatically scales the number of pods in a deployment, replication controller, or replica set based on that resource's CPU utilization. It requires a metrics source such as the Kubernetes Metrics Server to provide CPU usage data1. Cluster autoscaling is a feature of

Kubernetes that automatically adjusts the number of nodes in a cluster when pods fail or are rescheduled onto other nodes. It requires an integration with AWS Auto Scaling groups to manage the EC2 instances that join the cluster2. By using both horizontal pod autoscaling and cluster autoscaling, the solution can ensure that Amazon EKS scales in and out according to the workload.

NEW QUESTION 122

......